A window to a dark past

BY CARYN GILLOOLY

Washington, D.C.

"Never forget."

That's the motto of the U.S. Holocaust Memorial Museum here, which was set up to help visitors understand the magnitude of the atrocities committed during the scarifying days of the Third Reich. Included are displays of shoes from some of the more than six million people murdered and a boxcar used to transport victims to concentration camps.

But like other institutions, the museum can only share so much information via a walk-through tour. To supplement exhibits, the museum has developed an internal multimedia network accessible to visitors and staffers, and is extending its resources to the Internet

Holocaust Memorial Museum

community via an on-site World-Wide Web (WWW) server.

The network, which plays an inte-

IBM plans enterprise attack

BY MICHAEL COONEY

Raleigh, N.C.

IBM this month will rev up its multivendor management offerings with new NetWare control capabilities and software distribution tools, widen the reach of its peer network technology and refocus its high-speed networking initiatives.

The company will announce a new version of NetView Performance Monitor (NPM) that supports Novell, Inc. NetWare local networks as well as a more powerful version of NetView Distribution Manager/6000 — its software

distribution tool. IBM will also add Advanced Peer-to-Peer Networking (APPN) support to the 8250/8260 hub family, a move that could make it easier to build and manage enterprise nets.

What's more, IBM is expected to recast its AnyNet software, possibly combining it with High Performance Routing products to jump-start interest in these strategic technologies. In addition, IBM will give users some solid Asynchronous Transfer Mode (ATM) product delivery dates.

"IBM is articulating how it will tie

Listen up, IBM

Big Blue's big users sound off in our special issue on IBM networking. Also included: an inside look at SNA routing, insight into the evolution of APPN and a frank discussion with IBM's top net management executive. Page 47.

NEWSPAPER

AT&T plots to grab lucrative toll traffic in local loop raid

BY JOANIE WEXLER

AT&T has launched a quiet assault on the local toll market, using financial incentives to get customers to snub the local carriers and route toll calls over its own network.

As part of its plan for grabbing a big share of the $10 billion intra-LATA toll market now owned by local carriers, AT&T has begun offering custom discounts to users that reprogram telephone switches to bypass the local carrier, an AT&T spokesman confirmed.

In addition to price savings, users could benefit financially by lumping local toll calls with other AT&T network traffic to generate larger volume discounts, said Steve Sazegari, senior analyst with Ryan Hankin Kent, a research firm in South San Francisco.

The discounts are available to business users in a growing number of states that permit competition for toll calls that do not cross local access and transport area

boundaries (see graphic, page 86).

AT&T will work with customers of its Definity private branch exchanges to do any reprogramming needed to deliver local toll calls automatically into the AT&T network. AT&T will also sell autodialers to do the job for small users or those with PBXs too old to upgrade.

"This might be desirable because I wouldn't have to rely [solely] on the [local exchange carrier] anymore," mused Eric Shafer, communications analyst for Norwest Technical Services in Des Moines, Iowa. "Having all my calls go over one service would be less of a headache."

Currently, intra-LATA interexchange calls — those that travel outside a local calling area — are carried as toll calls by LECs, and the rates for such calls are regulated by the states.

In AT&T's new PBX scenario, the route decision-making would take place within the PBX, which would grab the intra-LATA calls and send them directly to the AT&T point of presence.

AT&T gives customers a ring

BY JOANIE WEXLER

Basking Ridge, N.J.

AT&T last week grabbed the reliability reins of the "last mile" in customer networks.

The carrier announced Accu-Ring, a service that allows customers to aggregate access to private-line, virtual net, 800, frame relay and other AT&T services over a self-healing fiber ring.

> **AT&T boosts reliability of Accunet app. Page 87.**

Accu-Ring is targeted squarely at AT&T's largest customers, as it requires a minimum aggregate access link of 45M bit/sec within a single local-access and transport area.

AT&T, which said it does not intend to get into the business of building fiber rings, has allied with several providers of redundant fiber networks to offer Accu-Ring. Allies include competitive access provider (CAP) Teleport Communications Group,

Mail management by the book

BY JIM DUFFY

Parsippany, N.J.

Struggling to manage your far-flung messaging network? Then look to the company that wrote the book on it.

AT&T has developed an automated management platform for its EasyLink Services messaging network that blends off-the-shelf products with custom software. The platform monitors and controls message transfer agents (MTA), mail gateways and switches using intelligent software agents and knowledge-based management applications that minimize human intervention.

According to an AT&T official, the technology could be applied to any large corporate network and can be developed and deployed in as little as eight months. In fact, the AT&T executive wrote a book on how to develop a management framework for messaging networks that will be released in September.

AT&T's framework relies on the Simple Network Management Protocol and some internally developed MIB, intelligent agent and element management software. This homegrown code is combined with off-the-shelf products to create an integrated, yet modular and scalable management environment.

"Our challenge was to come up with an enterprise management framework which doesn't lock us into any one vendor or one solution," said Raj Ananthanpillai, technical manager for AT&T Easy-

Raj Ananthanpillai of AT&T EasyLink Services

Briefs

ATM start-up to emerge. In two weeks, electrical connector giant AMP will introduce a subsidiary company to focus on development of workgroup Asynchronous Transfer Mode (ATM) products. The company is called Connectware, Inc., and its product line, CELLerity, will consist of ATM workstations, servers, network interface cards, stackable hubs and switches. Connectware is promising an aggressive pricing structure enabled by a highly integrated segmentation and reassembly chip.

Connectware: (214) 907-1093.

That secure feeling. RSA Data Security, Inc., which sells public-key cryptography software, and Enterprise Integration Technologies Corp., product manager for a $12 million Internet marketplace, are teaming up to sell tools that build secure Internet applications. Their joint venture, Terisa Systems, will market tools for creating World-Wide Web clients and servers that incorporate cryptography and digital signatures for ensuring the privacy and validity of transactions over the Internet. Shipping is scheduled for late 1994.

Terisa: (415) 617-1836.

Sequent's in the mail. Sequent Computer Systems, Inc., best known for its Unix-based on-line transaction processing systems, this week moves into enterprise messaging. The company is expected to announce a series of rollouts over the next few months of partnerships, as well as hardware and software tools for stitching together enterprise electronic mail backbones and for connecting corporations via electronic data interchange.

Rooms and desktops unite. AT&T is expected to announce Wednesday a strategy for tying desktop- and room-based multimedia systems together on an as-needed basis. AT&T's Business Visual Communications group and various AT&T transport business units are reportedly teaming with Lotus Development Corp., Novell, Inc., Xerox Corp. and Intel Corp. — as well as VideoServer, Inc. in Lexington, Mass. — on the effort.

MCI takes flight. This is a case of what goes around coming around. MCI Communications Corp. last week said it will acquire a minority stake in In-Flight Phone Corp., a company started by Jack Goeken, the man who founded the company in 1963 that grew into MCI.

In-Flight's system lets passengers make telephone calls, receive paging messages, send facsimiles, get stock quotes, view airport maps, play video games and order gifts.

A blow for the CAPs. The U.S. Court of Appeals for the District of Columbia Friday overturned the Federal Communications Commission's 1992 order that Bell operating companies must allow competitive access providers (CAP) to collocate switching equipment in BOC central offices. Within hours, one of the nation's two leading CAPs — Teleport Communications Group — filed a petition asking the FCC to revoke certain pricing flexibility it had granted BOCs as a quid pro quo for collocation unless the BOCs voluntarily continue to provide collocation. The other leading national CAP, MFS Communications Co. of Oak Brook, Ill., refused immediate comment.

Microsoft preview. Microsoft Corp. this week will unveil the company's new database strategy, revised in light of its recent split with Sybase, Inc. on SQL Server development. Jim Allchin, vice president for the business systems division at Microsoft, will discuss the next release of SQL Server for NT, among other things.

Correction: *Network World* last week ran what was believed to be a fictitious 800 number in a headline about new toll-free service options from Bell Atlantic Corp. That number is actually in service by another company and should not be used to reach Bell Atlantic. For more information on the Bell Atlantic services, call (201) 647-2279.

Contacts

Table of Contents

NEXT WEEK:
A report from the Summit

IT USER Summit

Representatives from key user groups will gather in Boston this week for the IT User Summit — Mapping the Future of Information Technology.

Network HELP desk

Network World *tracks down answers to your questions regarding products, services, technologies or disputes with vendors. Please submit questions to Dana Thorat at (800) 622-1108, via fax at (508) 820-3467, via the Internet at djt@world.std.com or via CompuServe at 73244,2673.*

Is PCMCIA supported by Banyan Systems, Inc.'s VINES 5.5X? If so, what particular brands of Ethernet boards are supported?

Bryan Miller, Richmond, Va.

Gary Wolfe, editor of public communications at Banyan, which is based in Westborough, Mass., replies:

Banyan has certified several PCMCIA Ethernet adapter cards to work with the VINES 5.X product line — IBM's Ethernet 10BT, Ethernet 10B2 and AT&T's WaveLAN/PCMCIA. In addition, IBM's Token-Ring credit card adapter, TRN16/4, is also supported by VINES 5.X.

Compatibility is based on how the Network Driver Interface Specification (NDIS) drivers in the adapter cards interoperate with the NDIS client drivers included in the VINES 5.X operating system.

Ethernet 10BT costs $225, and Ethernet 10B2 costs $240. For more information about the IBM products, call (800) 426-2255.

The AT&T WaveLAN/PCMCIA card costs $695. For more information about this product, call (800) 858-3435.

What is the availability of Asynchronous Transfer Mode (ATM) cards from Hewlett-Packard Co., Sun Microsystems Computer Corp., Silicon Graphics, Inc. and IBM?
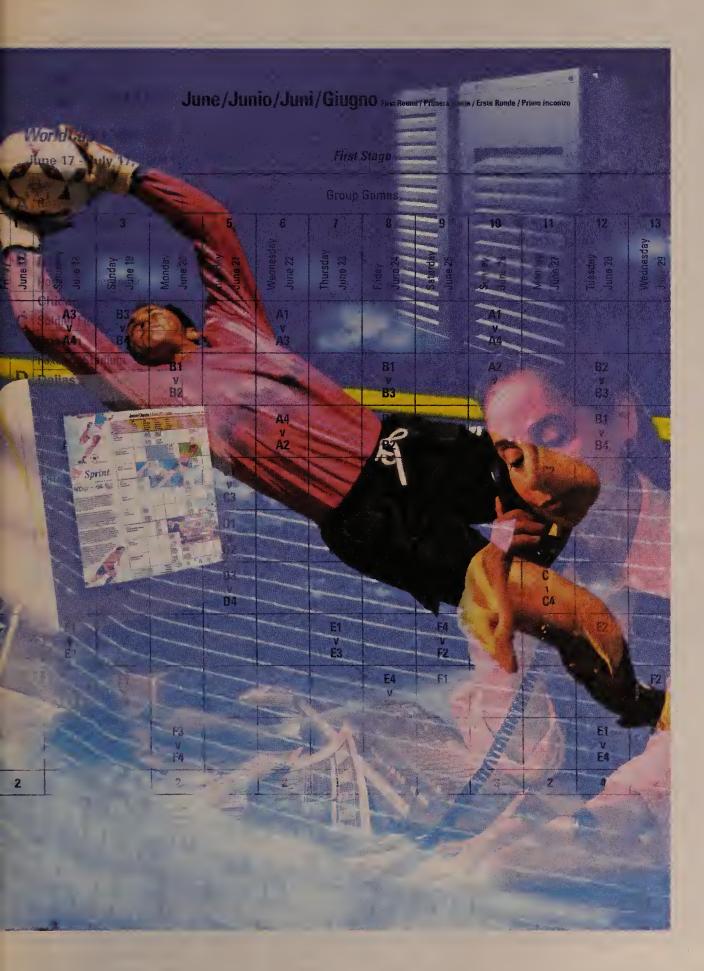
J.S., Menlo Park, Calif.

NW found that only two of these vendors currently provide ATM cards.

SMCC has announced two SunATM cards for SBus-based SMCC workstations. One supports multimode fiber networks, and the other supports unshielded twisted-pair Category 5 cabling. Both cards support speeds of 155M bit/sec. The multimode fiber network ATM card will be available in August and will cost $1,295. The ATM card that supports twisted-pair cabling will be available in September and will cost $995. For more information, call (800) 821-4643.

Router vendors bulk up remote office devices

BY SKIP MacASKILL

Internetworking vendors are reaching out to the farthest flung reaches of the corporate empire with new products and features that tie even the smallest sites into the enterprise.

Advanced Computer Communications (ACC) this week will roll out a new branch office router priced under $2,000, while 3Com Corp. last week took steps to fully integrate ISDN in its product line.

"In terms of unit volume, the remote access market is the fastest growing area of the overall router market," said Melinda Le Baron, an analyst at Gartner Group, Inc., a consultancy based in Stamford, Conn. "The technology has matured and price points have dropped low enough that companies are more willing to internetwork even the smallest branch offices."

ACC will unveil Danube, a scaled-down version of its Nile router that provides full routing and band-

When the remote office Danube is used in conjunction with our other routers in regional and corporate offices, we can offer enterprise routing solutions priced 40% to 50% lower than our competitors."

Gary Krall
Vice president of marketing
ACC

width management. The device supports a single 10Base-T or 10Base5 Ethernet interface and one wide-area network connection to frame relay, ISDN and X.25 networks for only $1,995.

Gary Krall, vice president of marketing at ACC,

said the company is tapping into the manufacturing resources of its parent company, Newbridge Networks, Inc., in order to produce a low-price product.

"The upside for the user is that for just under $2,000, they get a full-fledged router, not enhanced bridges or IP-only devices like our rivals are offering at that end of the market," Krall said.

Danube will support ACC's suite of bandwidth optimization features, including 4-to-1 data compression; Express Queuing, which automatically allocates bandwidth to sessions; and Dial-on-Demand, which will automatically dial up and tear down switched connections, reducing WAN costs.

"The significance is the inclusion of these bandwidth management features in a product priced under $2,000," Le Baron said. "ACC may be the only company that offers data compression over frame relay, which can help users [cut] WAN costs even more."

Danube can handle Internet Protocol, Internetwork Packet Exchange (IPX) and AppleTalk Phase I and II traffic, with Open Shortest Path First and Routing Information Protocol routing as well as transparent bridging. It is available now.

3COM AND ISDN

At last week's NetWorld+Interop show in Berlin, 3Com announced that it will fully integrate ISDN terminal adapter capabilities into its NetBuilder II, NetBuilder Remote Office and AccessBuilder routing line.

"By closely folding ISDN capabilities throughout our central, remote office and at-home office lines, users will be able to forge end-to-end ISDN links across the enterprise," said Janice Roberts, 3Com's vice president of marketing.

In the fourth quarter, 3Com will unveil an ISDN Primary Rate Interface (PRI) for its NetBuilder II backbone router, an integrated four-port Basic Rate Interface (BRI) module for its AccessBuilder router as well as an Industry Standard Architecture adapter card with a built-in BRI interface.

In 1995, 3Com will follow with integrated BRI and PRI modules for the NetBuilder II, a PRI module for AccessBuilder and a PCMCIA adapter card with BRI capabilities. "ISDN is still not deployed enough in the U.S. for these rollouts to make much of a difference," Le Baron said. "Integrated ISDN is in big demand in Europe, which is why 3Com probably made this announcement at the show in Berlin."

ⒸACC: (408) 366-9600; 3Com: (408) 764-5000.

Chipcom casts its dice in switch mgmt. game

BY SKIP MacASKILL

Southborough, Mass.

Chipcom Corp. this week will unveil a network management strategy and applications for dealing with a mix of shared and switched LAN segments in an enterprise net.

The hub maker said it will code-velop a suite of management and traffic analysis applications with Axon Networks, Inc. and Hewlett-Packard Co. that will be integrated into Chipcom's ONdemand Network Control System (NCS) net management package.

These new applications will complement those users already employ to handle shared networks by adding support for local nets based on Chipcom switching hubs and allowing users to manage both types of environments from the same management console.

Competition in the market for managing switched and shared LANs is heating up. In recent weeks, SynOptics Communications, Inc. has teamed with Network General Corp. and Hughes LAN Systems, Inc. has partnered with HP to address the issue.

John McConnell, president of McConnell Consulting, Inc. in Boulder, Colo., said the key to managing existing LAN internets and emerging switched LANs in a unified way is to install probes and analysis tools in a convenient place, then manage those environments from a single console — exactly what Chipcom plans to do.

"The hub is the most logical place to house those probes and kick off the needed applications, because in a switched environment, there's no longer a common piece of wire to drop a Sniffer onto," he said.

Chipcom's net management architecture consists of five layers — Domain, Device, Advanced Network Analysis, Logical and Virtual — representing a variety of monitoring, analysis and configuration capabilities.

Under the Domain layer, Chipcom and its partners will deliver software agents for all Chipcom devices that will compile a variety of device-specific information. The Device layer refers to a series of applications with graphical user interfaces (GUI) that will correlate all of the agent information.

At the Advanced Network Analysis layer, Chipcom will provide net and protocol analysis, fault isolation and problem determination functionality for Chipcom-based nets and other networks supporting Remote Monitoring (RMON) probes. Chipcom will accomplish this by deploying RMON technology from Axon in the form of two new LANsentry applications — Statistical Analyzer and Packet Decoder — which will be integrated with Chipcom's

ONdemand NCS.

The Axon analyzer application reviews both real-time and historical data, allowing users to establish a baseline and analyze trends for comparing current and past network performance. The decoder tool allows users to customize filters and create buffers for gathering and storing information on specific packet types and protocols.

Chipcom readies new mgmt. lineup		
Product	**Price**	**Availability**
LANconnect	$2,495	30-60 days
LANsentry Packet Decoder	$3,995	45-60 days
LANsentry Statistical Analyzer	$3,995	45-60 days
NetMetrix Internetwork Monitor for ONdemand	TBA	1Q 1995
ONdemand NCS (Unix version)	$4,995	30 days
ONdemand NCS (Windows - desktop version)	$2,995	30 days
ONdemand NCS (Windows - portable version)	$895	30 days

TBA = To be announced

"Providing RMON information is one thing, but analyzing it and breaking it down into useful tidbits is key," said Charlie Robbins, vice president of communications research at Aberdeen Group, Inc in Boston.

The Logical layer involves the integration of HP's existing NetMetrix Internetwork Monitor with Chipcom's ONdemand NCS. The HP application is designed to let users dynamically reassign end nodes and reconfigure nets via knowledge-based modeling and performance analysis technology.

The Virtual layer, a superset of the Logical layer, defines products that will allow net managers to control and monitor several virtual LANs that are spread across an enterprise.

"In a virtual net, you'll be able to group users together across the enterprise without worrying about address changes or router boundaries," said Phil Fulchino, director of enterprise systems at Chipcom.

In addition to the LANsentry applications, Chipcom will also roll out LANconnect, a new GUI-based application for managing internetworking modules in Chipcom's ONline and ONcore intelligent hubs. Based on technology from NetLabs, Inc., LANconnect integrates with ONdemand NCS and allows for configuration and monitoring of all router, bridge and switching modules in the hubs.

New versions of ONdemand NCS — Version 3.0 for Unix and Version 1.5 for Windows — will also debut. New features include support for Chipcom's newest hubs, the ONcore and ONsemble stackables, as well as the new management applications.

The Windows package will be available in a portable version, which operates as a stand-alone product, and a desktop version, which will come bundled with HP OpenView.

ⒸChipcom: (508) 460-8900.

Fed users urge radical power shift

BY ELLEN MESSMER

Washington, D.C.

As expected, a federal technology panel has called for an end to the Government Open Systems Interconnection Profile (GOSIP) procurement mandate and asked a little-known group to take the reins on interoperability efforts.

The National Institute of Standards and Technology, the agency with authority to set standards requirements, has already acquiesced to these demands that were contained in the Federal Internetworking Requirements Panel (FIRP) report issued last week.

The report asks the Government Information Technology Services (GITS) Working Group — just one part of the Information Infrastructure Task Force (IITF) set up at Vice President Al Gore's request — to coordinate government network-compatibility efforts without mandatory standards.

Considering the FIRP comprises eight government network managers, its request essentially means the new federal networking strategy will be based on grass-roots voluntary user support — something the GOSIP mandate never really had.

In its report, the FIRP admitted agencies often

FLYSIK

bought OSI products and didn't use them, or found ways to dodge the purchasing mandate.

The report claimed that OSI products are generally more expensive, less available and harder to use than comparable Transmission Control Protocol/In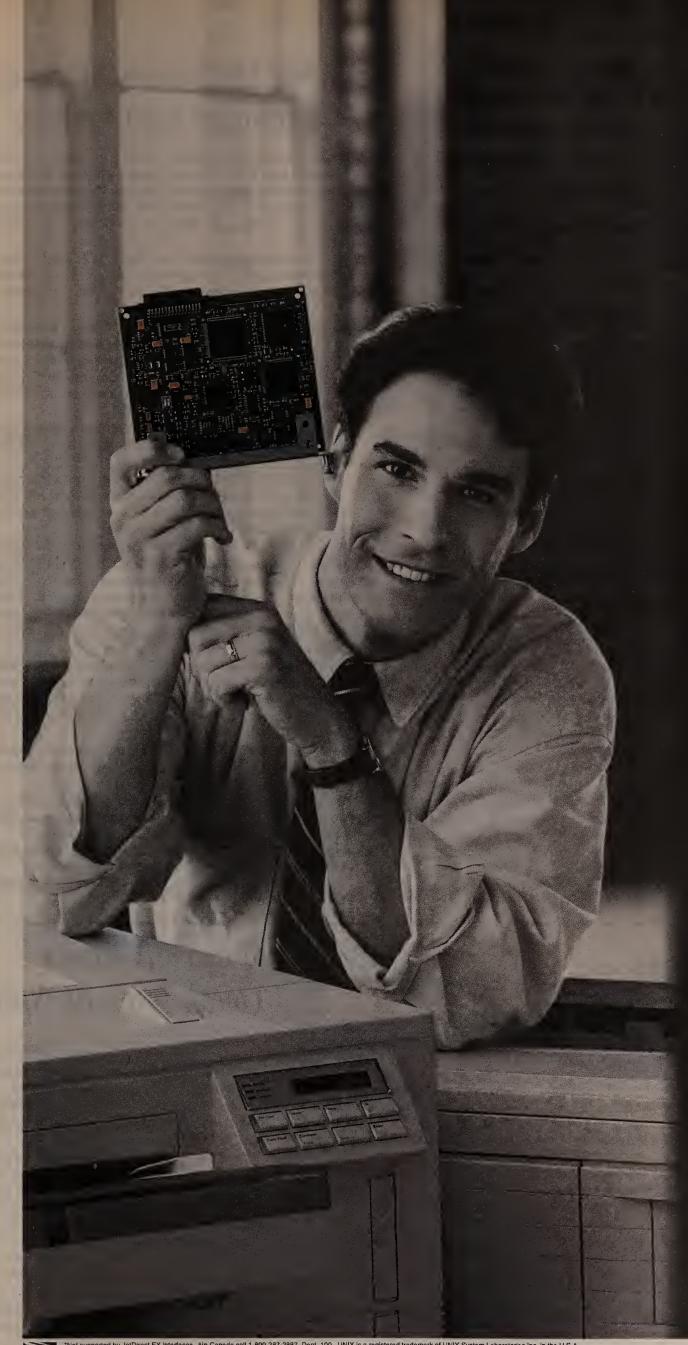ternet Protocol products. However, it concluded that no single protocol suite meets the full range of government requirements for data networking.

"[OSI] is an issue that's raged for some time," said Robert Woods, associate administrator for Federal Telecommunications System (FTS) 2000 at the General Services Administration. "If you set a standard and nobody comes, is it really a standard?"

The network managers who wrote the report have demanded an end to the five-year government policy backing OSI, instead proposing that "affinity groups" be set up to define network interoperability on their own terms. The groups would deal with subjects such as electronic data interchange and electronic mail.

Acknowledging the need for some coordination, the report asks the White House Office of Manage-

FCC lays out revamped PCS spectrum plan

BY DAVID ROHDE

Washington, D.C.

The FCC last week set the stage for personal communications services (PCS) auctions to soon take place by succumbing to industry demands for a spectrum allocation plan that places all the wideband licenses adjacent to one another.

But in clearing the way for auctions — now likely to be held in December, sources say — the Federal Communications Commission did not silence doubt about the ability of PCS to serve as anything more than a cellular look-alike.

"Initially, you're going to see a heavy concentration in cellular-type [voice] services," said Mark Golden, acting president of the Personal Communications Industry Association. "The new entrants will do something that looks like cellular at first because they need to build revenue."

As expected (NW, June 6, page 39), the new spectrum allocation plan approved by the FCC includes three 30-MHz channel blocks and three 10-MHz channel blocks. Two of the three 30-MHz channels are reserved for Major Trading Areas, a region that may include multiple cities or states. The other blocks are Basic Trading Areas, which include only one metropolitan area.

Every area will have up to six PCS providers, whereas there is a maximum of two cellular providers today. That increase is expected to promote far more competition and, presumably, lower prices.

And each of the six blocks of spectrum are actually divided into two separate bands, with one for transmitting and one for receiving.

All six blocks are now clustered in an area just below 2 GHz on the radio frequency spectrum. The old plan included a combination of 10-, 20- and 30-MHz channels scattered below and above 2 GHz.

Making the channels adjacent obviates the need for separate handsets for different PCS bands, reducing likely equipment costs by at least 25%, said Robert Pepper, director of the FCC's Office of Plans and Policy.

And by removing PCS licenses from the area above 2 GHz, which is heavily populated by microwave users, the agency reduced the cost of paying off those incumbent users by an estimated $1 billion, Pepper said.

Within a couple of years, PCS should offer much more bandwidth than existing wireless services, which generally max out at 19.2K bit/sec. The promise is that users will be able to send multimegabit file transfers, images and facsimiles over 56K bit/sec wireless links.

Nonetheless, some fear PCS will muddy the waters for users searching for a viable means of wireless data transport.

"It's a clean slate for digital voice, but from a data point of view, I wouldn't jump on it," said Gerhard Blendstrup, senior vice president of strategic services for the ADP Claims Solutions Group in San Ramon, Calif.

Blendstrup expressed fear that the advent of PCS would freeze user plans for deploying other wireless technologies, such as specialized mobile radio or Cellular Digital Packet Data.

"Trying to bet on any of these technologies drives me nuts," Blendstrup said.

The PCS auctions are also causing concern among cellular carriers, even though the FCC last week granted them a slight concession.

Like the old rules, the new rules limit incumbent cellular providers to bidding on the 10-MHz channels if there is more than a 10% overlap between the cellular territory and the PCS territory.

But the FCC agreed to raise the overlap threshold to 20% if the cellular carrier agrees to scale back to a 10% overlap over time. The Cellular Telecommunications Industry Association said it was "very disappointed at some of the details of the decision, such as the . . . rules to determine how cellular carriers can grow in areas adjacent to their markets."

The FCC still must address the issue of whether to set aside some spectrum for companies owned by women and minorities, and determine the auction procedure in general. Those issues are expected to be decided on June 29. The auction could take place as soon as four months after that, but industry sources said December is a more likely target date. **Z**

> "Trying to bet on any of these technologies drives me nuts," said Gerhard Blendstrup of ADP Claims Solutions Group.

Comments?

See "Contacts" box on page 2.

Deadline fast approaching for PCS number applications

Potential bidders for personal communications services (PCS) licenses may think they have plenty of time to prepare for next month's narrowband auctions and the wideband auctions now expected in December. But they actually have to get moving this week — to apply for telephone numbers for their services.

Friday is the deadline to apply for a Service Access Code — the middle three digits, or exchange portion, of a 10-digit phone number — under the new 500 Area Code reserved for PCS.

DIAL June 17 500

The June 17 deadline was set by Bell Communications Research, based here, after the Federal Communications Commission last month released the new area code for use.

And Bellcore, which administers the North American Numbering Plan, already has a backlog of applications for PCS numbers.

Those who miss the deadline won't be shut out of numbers if they win licenses, said Harry Young, chairman of the Wireless Interconnection and Numbering Committee of the Personal Communications Industry Association.

That's because, in addition to the 500 numbers, millions of new telephone numbers will open up next Jan. 1 through the long-awaited revision of the North American Numbering Plan.

But the 500 number is likely to become identified with wireless the way 800 has become identified with toll-free service and sales. Other PCS numbers "may not be as easily recognizable," Young said.

Ironically, for a service intended for mobility, the 500 numbers will not be "portable" among carriers the way 800 numbers now are. Each service provider that gains a 500 Service Access Code will obtain control of all 10,000 possible line numbers behind it.

And applicants for 500 Service Access Codes can request specific three-digit number sequences, although there is likely to be much overlap among these requests.

A subcommittee of the Industry Numbering Committee — the industry forum assigned to resolve telephone numbering issues — is looking into portability of 500 numbers.

Bellcore (North American Numbering Plan Administration): (201) 740-4596.

BY DAVID ROHDE

CyberSpeak: *Voices from the reader network*

Is the FCC's effort to auction off wireless bandwidth a good idea, or are we mortgaging our children's network future?

"The Federal Communications Commission is charged with maintaining the electromagnetic spectrum's frequencies for the benefit of the public, as a public trust.

"To auction off bandwidth removes portions of it from the public domain, where all segments of business can interact under the rules that govern electromagnetic emissions, and creates a proprietary environment that would serve to hinder the growth of the communications industry.

"At the very least, implementation of the auction policy would require significant legal work to establish the 'boundaries' of the 'electromagnetic real estate' and to maintain it, leading to increased governmental bureaucratic expenses.

"I oppose the FCC's auctioning proposal."

Tom Honles, Structural Engineering Subsection, Los Angeles Department of Water and Power

"The radio spectrum is a limited resource and as such is public domain and should be used to the benefit of the people. As with the rest of the spectrum, you must show just cause to be entitled to operate there.

"If a service of greater value to the public requires your space, then they are awarded it. This 'make-good-use-of-it-or-lose-it' philosophy keeps us in the forefront of communications technology and puts us as close to the edge of state-of-the-art as is feasible."

Tony Podrasky, network specialist, Convex Computer Corp., San Diego

"If the FCC wants to auction bandwidth, that assumes the FCC owns it. I think our children will be happier if it auctions leases to the space and doesn't lock the resources."

Nick Catanese, Tricom Business Systems, Willoughby, Ohio

NextWeek CyberSpeak Out!

IBM is betting its network future on APPN. Are you?

Responses due by 8 p.m. Thursday, June 16. You'll get a T-shirt if we print your response. Please include your name, title, company and site, as well as your mailing address.

INTERNET: network@world.std.com * COMPUSERVE: 73244,2673 * FAX: (508) 820-3467 * RAF HOTLINE: (800) 622-1108, EXT. 487

Microsoft tosses Touchdown at Notes
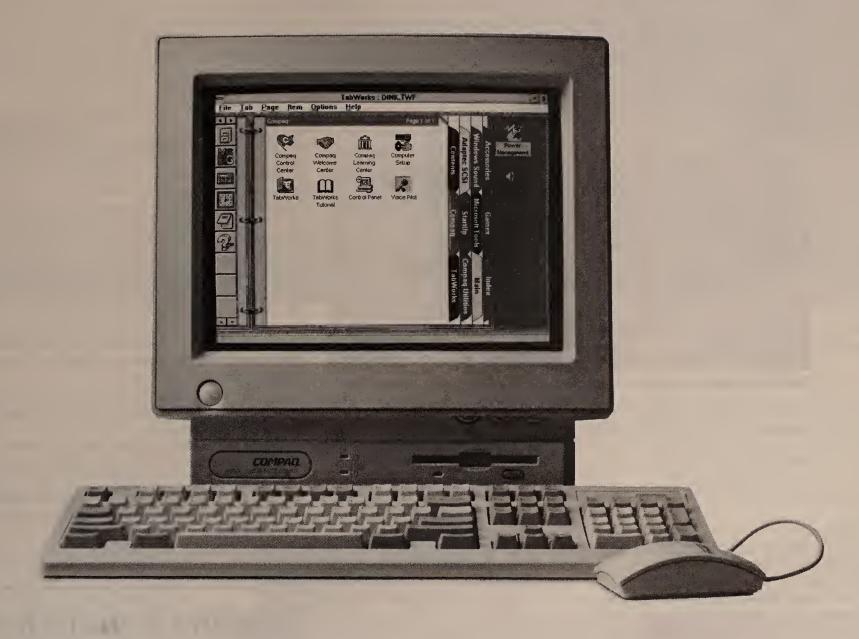
BY KEVIN FOGARTY

Redmond, Wash.

When Microsoft Corp. outlines its messaging strategy at a company user meeting here next week, it will position its upcoming client/server electronic mail product as a direct competitor to Lotus Development Corp.'s Notes.

While users and analysts are enthusiastic about the Microsoft Exchange, which is code-named Touchdown, they are having a tough time judging how its information sharing and work flow components will compare to Notes'. To date, Notes has faced little competition in these areas.

"As strictly a mail system, I think [Microsoft Exchange] will be quite compelling," said David Marshak, analyst at Patricia Seybold Group, Inc. in Boston. "It will be very interesting to see how comparable it will be, or if it will be able to leapfrog [the capabilities of] Notes."

Industry observers expect Microsoft Chairman Bill Gates to formally unveil Microsoft Exchange when the company demonstrates the product at the In-

IN A WORLD OF INSTANT COFFEE, INSTANT CREDIT AND INSTANT GRATIFICATION, WE NOW PRESENT INSTANT NETWORKING.

It could take a minute or two to get the Compaq ProLinea Net1 out of the box. But that's likely to be the most challenging part of your network installation.

Since it's completely self-contained, you simply put the ProLinea Net1 on a desk and plug it in. The NIC (Token Ring or Ethernet) and all the drivers are included, so it fits right into your network. It also opens easily for upgrades, so it fits into your future as well.

And to make the ProLinea Net1 fit into a business even better, for a limited time we're offering a special price on selected Token Ring models.*

To find out more, call 1-800-638-3258, ext. 435. We can direct you to a nearby authorized Compaq re-seller, or provide specifications on the ProLinea Net1.

It's the one PC that could solve quite a few of your networking headaches. Instantly. **COMPAQ**

Military girds nets for 21st century

BY ELLEN MESSMER

Washington, D.C.

When American and Allied combat forces stormed the beaches of Normandy 50 years ago, vital information about the battle was exchanged mainly over voice radios and face-to-face while commanders marked up strategy maps with grease pencils.

Military communications have progressed since, as coverage of the Persian Gulf War attests, but the U.S. is preparing for a future where commanders will plan strategy from computer-equipped armored vehicles, complete with electronic maps that display icons representing enemy and friendly forces in their near real-time positions on the battlefield.

The computer-based battle command system will be able to integrate reconnaissance data sent as images and messages over wireless networks from soldiers, field artillary and aircraft. From a mobile command post, officers will be able to quickly assemble a view of the battlefield and issue attack and movement orders to soldiers in combat.

These prototype electronic systems, to be used by ground soldiers and in tanks and other fighting vehicles, are now being tested at eight Battle Labs across the country (see graphic, this page) with the intent of preparing a Force 21 battle command for conflicts arising in the next century.

The battle command post of the future will be the Battle Command Vehicle, a large mobile armour-plated room that moves about on a tank-like chassis. Its nerve center will be General Dynamics Corp.'s Inter-Vehicle Information System (IVIS), a network of four workstations designated for Commander, Operations, Intelligence and Fire Support functions.

The Commander, Operations and Fire Support workstations will run software applications called Brigade Below Command and Control (B2/C2) to store and access text and graphic information about the battle, logistics support and combat-fire reports.

The USS George Washington: A floating model for military communications.

Labs testing 21st century battle systems

- **Battle Command Fort Gordon**
  Fort Gordon, Ga.
- **Battle Command Fort Huachuca**
  Fort Huachuca, Ariz.
- **Battle Command Fort Leavenworth**
  Fort Leavenworth, Kan.
- **Combat Support Service**
  Fort Lee, Va.
- **Depth & Simultaneous Attack**
  Fort Sill, Okla.
- **Dismounted Battle Space**
  Fort Benning, Ga.
- **Early Entry Lethality and Survivability**
  Fort Monroe, Va.
- **Mounted Battle Space**
  Fort Knox, Ky.

SOURCE: BATTLE LAB INTEGRATION & TECHNOLOGY DIRECTORATE, FORT MONROE, VA.

The Intelligence workstation will process target-sensor data passed over radio frequency networks to create a picture of enemy and terrain obstacles that can be shared with fighting forces using hand-held devices or laptops with B2/C2 software.

"We've field-tested it with soldiers," said Army Capt. Gerry Spragg, adding that the computer-based systems work reasonably well in combat exercises, but more development work still needs to be done.

As planned, IVIS data will initially be transmitted over the same radio network the military uses for voice traffic in combat situations — the Single Channel Ground and Airborne System (SINGARS) — built by ITT Corp. more than five years ago.

Sgt. Ricardo Castillo, who has tested Force 21 equipment at Fort Knox, noted that SINGARS was recently shown to work well for digital data at the platoon level, and tests are now starting at the larger brigade level.

IN THE HEAT OF BATTLE

But whether the radio net can simultaneously handle large amounts of voice and data traffic in the heat of battle is still being questioned. "That's the primary concern I have in this project," said Daniel O'Neil, Department of Defense electronics engineer for electromagnetic compatibility and analysis.

It will probably be several years before Force 21 ground soldiers cover terrain while carrying messaging devices and wearing helmets equipped with video cameras for transmitting images, as is the plan.

But the military continues to march on toward new technology.

Last week, the U.S. Navy announced that the aircraft carrier USS George Washington has become the first Navy vessel to make use of a new commercial broadband satellite service called the Maritime Telecommunications Network provided by Crescom Transmission Services, Inc. and SEATEL, Inc., a maker of antenna systems.

Under a Federal Communications Commission experimental license, Crescom is providing ocean-going vessels with the T-1 satellite links long enjoyed in ground systems. COMSAT World Systems, the U.S. partner in the International Telecommunications Satellite Organization, supported the Defense Department in convincing Congress that a commercial system, rather than the Military satellite system, could meet some defense needs.

Lt. Cmdr. John Hearing of the Naval Space Systems Command said the satellite system offers links to various land-based sites. As a result, services aboard the USS George Washington have been expanded to include telemedicine applications, image visualization, videoconferencing, remote database access and plain old telephone service for the crew.

Hearing last week said it is likely that most Navy vessels will use the T-1 service as a supplement to military satellite service in the future.

The USS George Washington, which has been testing the commercial service for six months, last week reached Normandy for the D-Day commemoration.

If the U.S. military was forced to cross the channel again to free Europe, it would be done differently today, mused one military source.

"Once we knew where the enemy was, we'd avoid them but cut off the supply lines and escape path," Hearing said.

"What we're trying to preclude is heavy-use bombs in the future. We want to use electronic warfare to zap his antenna and dominate the communications spectrum so we can talk, but they can't," he added. ∎

GTE tackles ATM packet loss problem

GTE Government Systems Corp. last week unveiled an enhancement that enables its Asynchronous Transfer Mode (ATM) switch to keep packet loss to a minimum, even when handling a mix of traffic that includes bursty TCP/IP data.

The vendor added a larger queue for variable bit rate (VBR) traffic to its Secure Prioritized ATM Network (SPANet) switch. The queue holds as many as 8,000 53-byte ATM cells at a time while constant-bit-rate traffic, such as voice and broadcast-quality video, moves directly through on a priority basis.

The deep VBR queue is designed to prevent packet discard, espcially for bursty Transmission Control Protocol traffic that is forced to compete with constant-bit-rate transmissions.

"That's been one of the really key issues that's been hounding ATM for a while," said John Morency, principal at Strategic Networks Consulting, Inc. in Rockland, Mass.

For example, a top MCI Communications Corp. executive recently said the carrier has found that buffer sizes below 5,000 cells cannot reliably provide 90% throughput of TCP/Internet Protocol traffic sent over ATM (*NW*, April 18, page 4).

Designed for military use, SPANet "is a rather expensive switch because it has security features that, for most users, are rather superfluous," said Lawrence Gasman, president of Communications Industry Researchers, Inc., a broadband services market research firm in Charlottesville, Va.

But Northern Telecom, Inc. is privately labelling essentially the same switch as the Magellan Gateway under an agreement with GTE Government Systems, according to Ken Napier, marketing director for broadband systems for GTE.

"Presumably, if [GTE Government Systems] they could upgrade SPANet to give it deep buffers, [Northern Telecom] could do the same thing for Gateway," Gasman said. "And I think that would be rather attractive."

MCI, which has shied away from committing to a national ATM service, reportedly has been testing the Magellan Gateway, a carrier access switch.

GTE Government Systems' move puts other ATM switch makers a step behind in providing predictable multimedia functionality, Morency said.

"There are lot of vendors out there who haven't done a lot with buffering," he said. "They've been concentrating on the base functionality."

BY DAVID ROHDE

Vendors look to link Notes with other apps, databases

BY ADAM GAFFIN

Boston

Percussion Software, Inc. last week introduced tools for linking Lotus Development Corp. Notes applications and databases with external databases.

The announcement comes as Lotus itself readies a formal launch for its Notes Visual Programmer (ViP) application development tool, also designed for stitching together Notes and external applications. ViP is set to ship by July, Lotus said.

Recent months have seen a spate of Notes tool announcements as Lotus and other vendors rush to address a common complaint of Notes administrators and developers: The Notes application program interface (API) is too complex to use for integrating the groupware product with other databases and applications.

Barry Reynolds, Percussion Soft-

ware president, said Notrix encapsulates the APIs for manipulating Notes databases within simple programming extensions to REXX, a high-level programming language developed by IBM that comes bundled with OS/2.

This lets developers quickly build server-based applications that can not only retrieve data from Notes, but also add or modify information, he said.

Reynolds said a key difference be-

■ CORRECTION ■

A story in the May 30 issue listed an incorrect title for John Walsh of Compression Labs, Inc. He holds a dual role as senior vice president/general manager, Multimedia & Personal Video Products and senior vice president, Corporate Planning & Strategy.

Grand Junction expands line with backbone device

Switch will act as aggregator of workgroup traffic.

BY SKIP MacASKILL

Fremont, Calif.

Grand Junction Networks, Inc. is aiming to help users build fast Ethernet backbone networks with the introduction this week of a 100M bit/sec Ethernet switch that links other workgroup switches.

The FastSwitch 10/100 Aggregator (AG) will support dedicated 10M bit/sec Ethernet links to the desktop, as well as switched and shared 100M bit/sec — or fast Ethernet — connections to other switches or servers.

The key difference between the AG and Grand Junction's existing FastSwitch 10/100 is an integrated fast Ethernet repeater, which divides 100M bit/sec of bandwidth between four ports. The repeater can be used to link other AGs or FastSwitch 10/100s, or provide access to a centralized server farm.

HESITANT ANALYSTS

While users praised this new capability, analysts were not as enthusiastic, saying the repeater could be quickly overtaxed in some environments.

The AG, which has a 1G bit/sec backplane and a 220M bit/sec forwarding rate, in addition features 25 switched 10M bit/sec Ethernet ports for high-end workstations or personal computers.

Also included is a switched 100M bit/sec fast Ethernet port that can connect the AG to a fast Ethernet backbone or a heavily accessed server.

"The AG is positioned as a low-cost alternative to a high-end switching hub, which carries a significant upfront investment," said Howard Charney, president and chief executive officer at Grand Junction.

"The AG allows users to add switches as needed and cascade them together," he added.

According to Charney, the AG offers users the best of two worlds. While classified as a workgroup switch, it also fills the role of a network switch by aggregating traffic from other switches.

Workgroup switches, including both the

> "Our old Ethernet network would drop users two or three times a day, but Grand Junction's switch technology put an end to that," said Fran Glasco, an AG beta user.

AG and FastSwitch 10/100, support a single media access control (MAC) address at each switched 10M bit/sec port and are designed to provide dedicated Ethernet links to high-performance workstations.
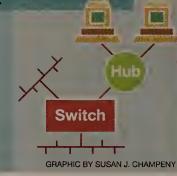
Workgroup vs. network switches

Workgroup switch:
Supports one MAC address at each port and provides dedicated 10M bit/sec of Ethernet bandwidth to high-end workstations.

Network switch:
Supports several thousand MAC addresses and is a low-cost router alternative that provides connectivity between shared Ethernet LANs.

GRAPHIC BY SUSAN J. CHAMPENY

On the other hand, a network switch typically supports thousands of MAC addresses. In addition, it can be a lower cost routing alternative for interconnecting standard Ethernet hubs.

Mary Petrosky, an analyst with The Burton Group, a consultancy based in Salt Lake City, did not think the AG offered a good method to interconnect other switches. "That four-port repeater shares a total of 100M bit/sec, which has the possibility of being gobbled up quickly in larger environments," she said. "The better approach would have been to provide dedicated 100M bit/sec to each of those ports."

SATISFIED USERS

Users, however, were pleased with the AG's new capabilities. Fran Glasco, an AG beta user and MIS director at Adolph Zimmerman Organization, Inc., a materials handling engineering firm in Memphis, Tenn., said the AG has cut downtime and boosted bandwidth.

"Our old Ethernet network would drop users two or three times a day, but Grand Junction's switch technology put an end to that," he said.

"The AG lets me establish a 100M bit/sec highway between my servers and the switch, as well as interconnect the four 10/100s we have on site. CAD files on the old 10Base2 system took 23 minutes to transfer, but with AG, that time has been reduced to seven minutes," he added.

The price was also right for Glasco. "The investment in the AG worked out to about $350 per user, and, if it meets our needs for the three years I'm expecting, it will turn out to be a very cost-effective investment," he said.

Available now, the AG costs $8,950.
©Grand Junction: (800) 747-3278.

Comments?

See "Contacts" box on page 2.

Oracle launches upgraded client/server applications

Includes 4 modules, support for distributed databases.

BY BARB COLE

Redwood Shores, Calif.

Oracle Corp. last week announced a new version of its client/server applications suite, including four new modules, improved manufacturing applications and support for distributed databases.

Oracle Cooperative Applications Release 10, a set of more than 25 applications, includes new cost management, sales and marketing, project costing and project billing modules.

Release 10 already includes Oracle Manufacturing, Oracle Financials, Oracle Human Resources and Oracle Project.

The applications are designed to run on AT&T Global Information Solutions MP RAS, Hewlett-Packard Co. HP-UX, IBM RISC System/6000, Sequent Computer Systems, Inc. DYNIX/ptx, SunSoft, Inc. Solaris, and The Santa Cruz Operation, Inc. SCO Unix clients and servers.

Also, Release 10 supports DOS and Windows clients, and support for an additional 20 server platforms will follow in 90 days.

Release 10 requires an Oracle database server, but it may access other data sources through Oracle's SQLNet middleware.

In addition to the new applications, Oracle has enhanced its manufacturing applications in Release 10 to support multiple languages and has more tightly integrated them with order entry applications.

The manufacturing modules, which were previously only available in English, are now available in 20 languages. Multicurrency support has also been improved, Oracle said.

WORTH THE WAIT

"This release has been long-awaited, but appears to be worth it because of improved functionality in the manufacturing modules," said Erik Keller, a vice president at Gartner Group, Inc., a market research firm in Stamford, Conn.

Ron Hawkins, director of computer operations at Millipore of Bedford, Mass., said there is better support for assemble-to-order manufacturing, which involves making products as orders come in.

In Release 10, users will be able to take customer orders at one location and check for inventory across their company's manufacturing operation, Oracle said.

DISTRIBUTING APPLICATIONS

Oracle also has added distributed concurrent manager software that lets network administrators define the number of application servers and the processes that run on each. These processors may be on one or more network servers.

"You may allocate resources by program, by user name or by application," said Don Klaiss, Oracle's vice president of manufacturing products.

In addition, each module in the suite includes an internal monitor that automatically reroutes work to another processor if there is a bottleneck.

Users said that Oracle has taken a big step toward distributed applications with its Release 10, but further enhancements are still needed.

Release 10 is aimed at companies that are downsizing business applications from mainframes to servers.

Each Release 10 application has a base price of $10,000 to $30,000, plus an additional $1,000 to $5,000 per user. The software will ship in July.

©Oracle: (415) 506-7000.

INTERNET tips

BY ADAM GAFFIN

One in a series of tips on Internet-based information services.

Government electronic commerce

The federal Electronic Commerce Acquisition Team is developing standards for how the government will do business electronically. A copy of its draft report can be obtained via anonymous FTP.

To access:
Use FTP to connect to ftp ds.internic.net. Log on as "anonymous," using "guest" as the password. Use the cd command to switch to the pub/ecat.library/march.architecture directory. Then use cd to switch to either the ascii directory for ASCII versions of chapters or to the postscript directory for Postscript versions (which include charts).

A NEW COMPANY. WITH 15 YEARS OF GLOBAL EXPERIENCE.

Every day the job of enterprise networking gets more complex. And every day for 15 years, we've been working with hundreds of the world's most successful companies to make things simpler.

Now we've simplified our name—from Ungermann-Bass to UB Networks—and streamlined our organization. To help you design and implement complex networking solutions, step by step. And provide you better access to information.

Anytime, anywhere.

SCALABLE, FLEXIBLE PRODUCTS.

We offer a wide range of connectivity, from Access/One intelligent hubs for enterprise solutions to the Access/Workgroup stackables—low-cost, high-performance networking for workgroups and remote offices. And our unique modular ATM Anywhere strategy offers the most cost-effective ATM connectivity available.

All our products are built for a long lifecycle to protect your investment. Instead of making drastic, expensive changes every few years, you can evolve into new technologies on a pay-as-you-grow basis.

Solve the networking puzzle yourself, or fold A to B and see how much easier it is to work with us.

A

START

STRATEGIC PARTNERSHIPS.

Your networking solution may include needs we can't supply alone. So we've assembled a team of strategic partners that's a who's-who of leading-edge technology. Tandem, Newbridge, NetWorth, Wellfleet, ACC, Fujitsu and LightStream, and more. Of course, we provide single-vendor support for solutions we specify.

Find out how we can help simplify your networking needs. Call for our free guide, "How To Solve Complex Networking Problems" and receive a poster of this ad.

We're ready, willing, and able to make things simpler for you.

B

ACCESS/ANYWHERE℠ from UB Networks

Access/Workgroups™
Stackable hubs and management tools for Ethernet and Token Ring workgroups with Access/Stax™ and Access/TRax™.

Access/One®
Full-featured manageable connectivity options for the entire enterprise using Token Ring, FDDI, ATM and Ethernet with VNA™ supporting virtual networks.

Access/Management™
The NetAssistant™ for Windows,™ NetDirector™ for UNIX® and OS/2™ products offer the network management features you need, priced for networks of any size.

Access/ATM™
Access/ATM products distribute high-bandwidth capabilities when and where you need them—ATM/Anywhere℠ for high-performance workgroup or LAN-to-WAN connectivity at the right price.

Access/Service™
Expert service, integration, consulting, and training available worldwide. And we can support your network 24 hours a day, 7 days a week.

For information call 1-800-777-4LAN.

FINISH

STARTING
WITH
OUR NEW
NAME.
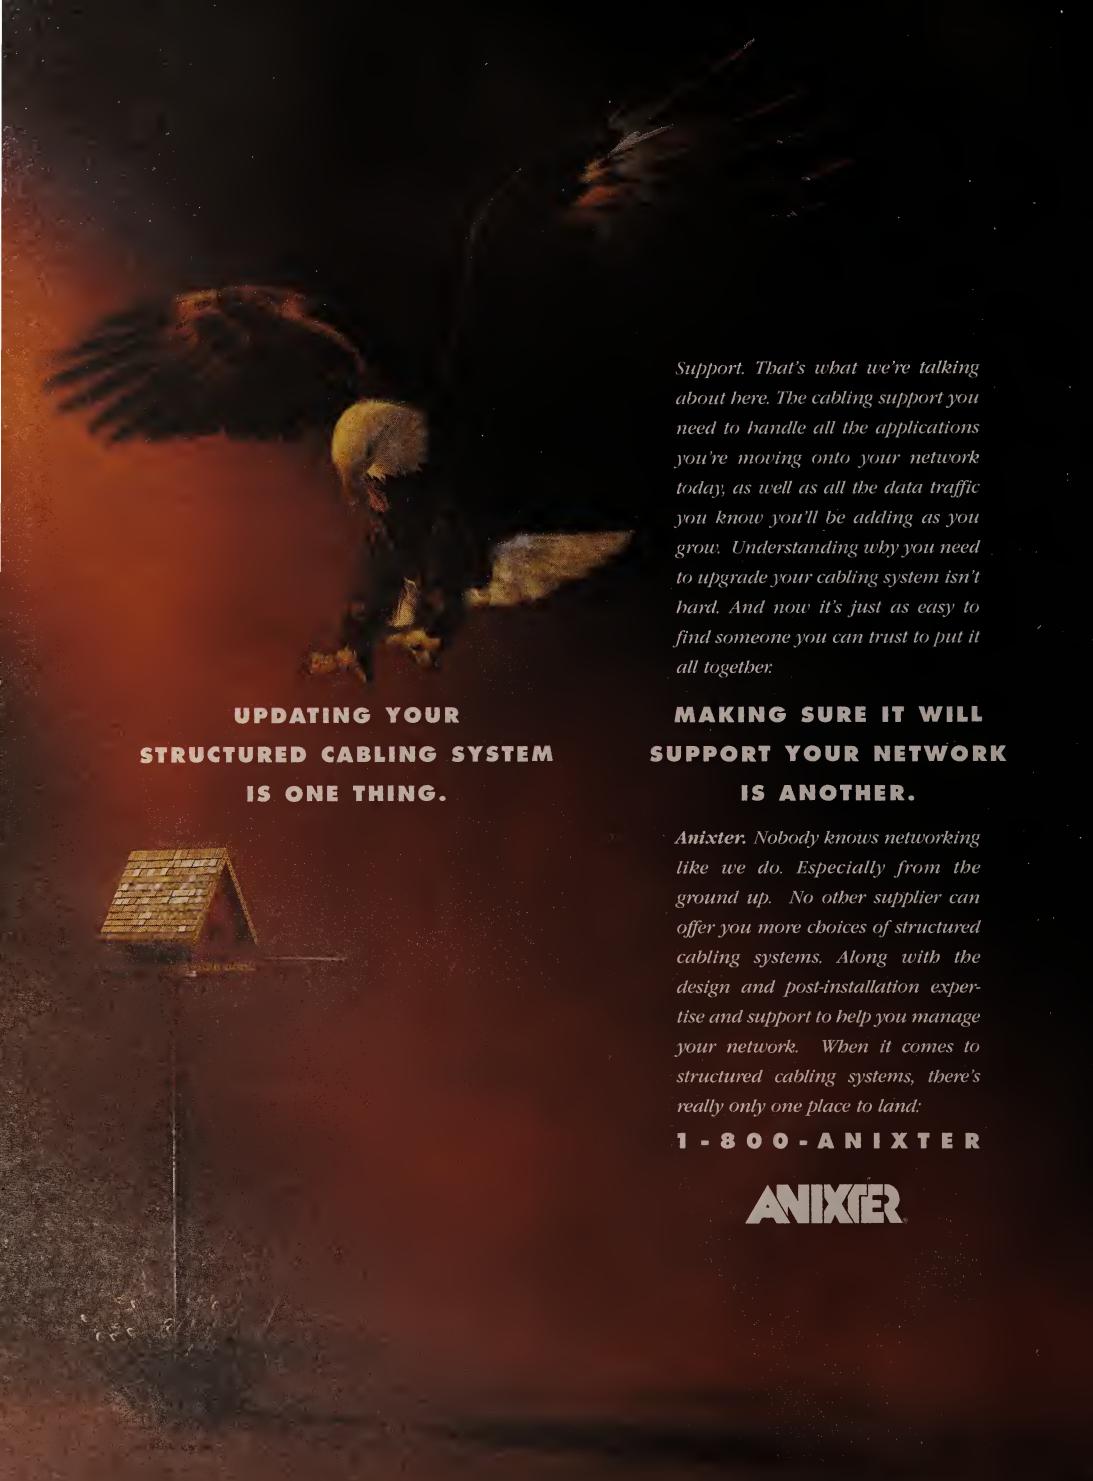
UB Networks
Formerly Ungermann-Bass

(408) 496-0111 Fax (408) 970-7337

Our Experience Makes It Simple.

MULTIPROTOCOL NETWORKING

IDEA opens up Concert doors to TCP/IP, Ethernet

BY MICHAEL COONEY

Billerica, Mass.

IDEA will this week roll out a new suite of software modules for its Concert communications controller that is designed to let SNA users more easily communicate with the TCP/IP world.

The new products will add support for Transmission Control Protocol/Internet Protocol Telnet and tn3270 sessions that will let users mix and match their TCP/IP and Systems Network Architecture devices on a single backbone. IDEA will also deliver a module that lets users link personal computers to Application System/400 hosts over Ethernet local-area nets.

"We are looking to provide enterprisewide access [using] a variety of networking protocols while providing concurrent access to SNA applications and new client/server applications on Unix hosts," said Rohit Mehra, product manager for IDEA's Concert line.

The Concert family includes five models that offer a variety of configurations, from the smallest Model 10100, which supports a maximum of eight 3270 or 5250 sessions, to the high-end Model 10600, which supports 128 3270, 84 5250 and 32 asynchronous sessions. Only the Model 10100 does not support the new modules.

The first new module, Concert Telnet Client software, resides on the Concert controller and gives downstream 3270 devices simultaneous access to TCP/IP and SNA hosts. Depending on model and configuration, Concert controllers can support from 32 to 128 sessions with either type of host. Each terminal can have as many as six sessions running at any time.

IDEA also will roll out Concert Telnet Async Server, a module that will let any existing asynchronous devices or Telnet clients simultaneously access IBM mainframes supporting TCP/IP applications and asynchronous host resources on TCP/IP nets.

The key feature with Telnet Async Server is that it acts like a gateway and eliminates the need to run TCP/IP on the mainframe for Telnet connectivity. With this new module, Concert converts the ASCII data streams from Telnet clients into the mainframe's EBCDIC character code and ships them to the destination application.

"Without having to run TCP/IP on the mainframe, users will save thousands of dollars in software charges," Mehra said. "Plus, none of the TCP/IP processing will be done on the mainframe, saving valuable cycles."

The Telnet Async Server will support as many as 16 terminals and as many as 64 sessions with the mainframe.

A new Concert tn3270 Server module will let tn3270 clients have access to SNA hosts over a TCP/IP net. Concert outfitted with the tn3270 Server module can support a minimum of eight users running 32 sessions to a maximum of 128 users running 500 sessions with the mainframe, depending on configuration.

Finally, IDEA will announce a Concert SNA-over-Ethernet module that will let 5250 terminals access AS/400s over 802.2-compliant Ethernet LANs.

Analysts said the TCP/IP enhancements will help IDEA stay even with and even push ahead of competitors such as IBM. IBM's 3174 family supports Ethernet, Telnet and tn3270 traffic, but only if users buy new controllers that support the features. Concert users can add these new modules to existing units.

IBM's AS/400 controller, the 5494, does not support Ethernet at all.

IDEA's tn3270 Server software costs from $995 for a 32-user version to $8,345 for 500 users. Telnet Client software lists for $495. The Telnet Async Server costs $795, and the SNA-over-Ethernet module runs $495. All products will be available by the end of July.

©IDEA: (508) 670-8512.

A market snapshot of controllers

	Shipments		Installed bases	
	1992	1993	1992	1993
IBM	13,520	7,500	225,250	201,213
Memorex-Telex	10,145	7,858	51,233	51,406
IDEA	1,564	1,238	5,207	4,986
Harris Adacom Corp.	1,160	757	3,899	3,886

Communications controller shipments fell 32% between 1992 and 1993, and IDC predicts the growth rate will fall an additional 29% between 1993 and 1998.

SOURCE: INTERNATIONAL DATA CORP., FRAMINGHAM, MASS.
GRAPHIC BY SUSAN J. CHAMPENY

HP rolls out new version of systems mgmt. software

OperationsCenter 1.1 offers more integrated mgmt.

BY JIM DUFFY

Palo Alto, Calif.

Hewlett-Packard Co. this week will unveil a new version of its systems management software that features enhanced multivendor event handling and application integration.

With the new product, users will be able to extend their management reach over a greater number of networked devices and more easily share data among management applications.

The new systems management software is called OperationsCenter Version 1.1. It runs on HP's Unix-based OpenView 3.3 platform and works with intelligent software agents to monitor system performance, distribute software, perform unattended backup and other functions.

Those software agents, meanwhile, can now work with AT&T System 3000 and SunSoft Solaris 2.3 platforms. They already supported IBM RISC System/6000, Sun Microsystems, Inc. SPARCstation, Bull HN Information Systems, Inc.'s DPX/20, and HP's 9000 Series and 3000 Series systems.

Version 1.1 features a number of new capabilities for OperationsCenter. Foremost is the software's ability to manage networked systems that run agent software other than the HP Intelligent Agent packages with which it usually works.

OperationsCenter can receive and display alerts and messages from HP and third-party Open-View applications used to manage subnets of an enterprise network, such as Systems Network Architecture, Novell, Inc. NetWare and DECnet networks. Since they are all built on the OpenView platform already, they have their own graphical map and event collection capabilities.

HP has now instrumented these applications so that OperationsCenter users can now click on an icon to pull down SNA, NetWare or DECnet maps and consolidate events from these environments into OperationsCenter's Message Browser database.

Message Browser is a log of active and inactive system events.

Version 1.1 also enables users to customize their layout of on-screen windows and map views. With this func-tion, systems administrators can set up multiple systems and deploy multiple operators, each with control over a specific piece of the enterprise network (see graphic).

Other features of OperationsCenter 1.1 are the ability to broadcast configuration information to hundreds of distributed agents; analyze a message log while resolving a problem to find ways to filter out irrelevant messages; and automatically respond to system-generated events at startup, such as remaining disk space, kernel messages and log files of applications.
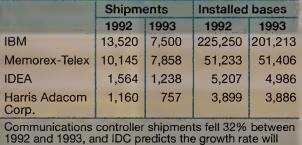
Analysts said they were impressed with HP's progress in making systems management more integrated.

"They're finally pulling it together a little bit more," said Charles Robbins, vice president of data communications research at Aberdeen Group, Inc. in Boston. "I like the level of integration I saw from how the screens are going to work. It all had a very uniform interface."

OperationsCenter will be available in August for a cost of $32,000.

Users are interested in OperationsCenter 1.1 but are a little concerned about its price.

"I was waiting to get a copy of that release and see if I could cost-justify it versus other tools that are just coming out," said Frank Belland, senior com-munications consultant at Martin Marietta Corp. in Orlando, Fla. "We are getting ready to look at another vendor's product that's really super cheap because some of the concerns here are the cost of technology."

The "super cheap" product is Network Partners, Inc.'s Trapper. It only costs $495 per server, Belland said.

©HP: (800) 752-0900.

Multiple personalities

HP's OperationsCenter 1.1 allows systems administrators to delegate different responsibilities to individual network managers. It also:

▶ Handles events from nodes without HP agents.

▶ Provides direct access to submaps from non-OperationsCenter applications.

▶ Offers a central display of all enterprise problems.

▶ Fosters process-oriented problem resolution.

Individual mangement domains or...

SNA   NetWare   DECnet

...all domains from a single console.

GRAPHIC BY SUSAN J. CHAMPENY

BRIEFS

For a limited time, users can purchase any **IBM 3174** communications controller model or feature and any **3299 Terminal Multiplexer** at a 40% percent **discount.** Orders for the equipment must be placed before Dec. 30. The promotion is only available through IBM Direct service at (800) 426-2255.

Hewlett-Packard Co. and **Novadigm, Inc.** last week said they will port Novadigm's object-oriented systems management framework to HP's HP-UX systems.

Under the agreement, both companies will cooperatively market an HP-UX version of Novadigm's **Enterprise Desktop Manager (EDM)** when it becomes available at the end of the third quarter. EDM performs configuration, security, change and asset management, and software distribution in enterprisewide client/server environments.

Client versions of EDM for HP-UX will cost $100, and server versions will cost $20,000.

HP: (800) 752-0900; Novadigm: (201) 512-1000.

Apertus Technologies, Inc. last week announced the availability of software that links Digital Equipment Corp. users to IBM hosts.

Called **Local Area Transport (LAT) Transport Services,** the software permits users of Digital's VAX terminals to log on directly to IBM mainframes and Application System/400 host computers via native mode 3270 and 5250 emulation. This is said to eliminate the requirement to log on to both the VAX and IBM systems and reduces the number of VAX user licenses required. LAT Transport Services is available now for $8,000.

Apertus: (612) 828-0300.

Router vendors bolster lines with new additions, capabilities

BY SKIP MacASKILL

Cisco Systems, Inc. last week continued to focus on easing router configuration by announcing a program that will let even network amateurs design, install and configure small to midsize router-based internets.

In separate actions, Cayman Systems, Inc. this week will introduce wide-area network expansion modules for its i Series of routers, while Sumitomo Electric U.S.A., Inc. will roll out two new Fiber Distributed Data Interface bridge/routers.

With new rivals such as hub and switch makers competing more directly in enterprise net areas traditionally controlled by routers — including the campus backbone and remote office — router vendors have been fighting back by adding functionality and making products easier to use.

Cisco's new Point and Click Internetworking program is a good example. It is a hardware/software package that was created for users that need to interconnect as many as 15 sites. The program bundles two or more Cisco 2500 entry-level routers with a suite of Windows-based software that steps the user through the configuration of multiple 2500s at various sites.

For small internets that do not require the range and functionality of a Cisco 4000 or 7000, users can stack multiple 2500s together at the central site and deploy individual 2500s in remote locations. The four members of the 2500 line offer a single Ethernet or token-ring port and varying combinations of synchronous serial, asynchronous and Integrated Services Digital Network Basic Rate Interface ports.

The software component ships on three CD-ROMs, the first of which is an internetworking tutorial that discusses design and installation issues, as well as descriptions of wide-area devices and services that complement routers. Another CD-ROM includes Cisco's Configuration Builder applications that allow users to configure multiple Cisco routers simultaneously and a program that steps the user through installation. Trivial File Transfer Protocol server and Telnet client software are also included to let users dial in to and configure remote routers.
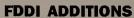
The final CD-ROM contains all needed documentation and allows users to search selected topics through the use of key words or phrases.

Point and Click Internetworking will be available through Cisco's European resellers this summer and domestically by year end. Pricing starts at around $7,000 for two 2500s and the CD-ROM pack.

Cisco held the lion's share of the 1993 worldwide FDDI router market based on ports shipped, according to Rising Star Research, a market research firm in Lakewood, Colo.

Proteon
Other
Cisco 54%
6%
5%
7%
8%
NSC 20%
Wellfleet
3Com Corp.

FDDI ADDITIONS

Sumitomo's new SumiNet-3500ME/4 and SumiNet-3500MT/2 FDDI bridge/routers allow as many as four Ethernet and two token-ring segments, respectively, to connect to an FDDI backbone via a dual-attached interface.

Based on Advanced Micro Devices, Inc.'s 29030 processor, the devices add to Sumitomo's existing 3500H series, which only provides FDDI bridging.

Network protocols supported include Internet Protocol, Internetwork Packet Exchange (IPX), AppleTalk II, DECnet Phase IV, Xerox Network Systems and Open Systems Interconnection. The Routing Information Protocol and Open Shortest Path First routing protocols are supported along with the Spanning Tree, source route and source route transparent bridging algorithms.

Available now, the Ethernet version costs $10,900 and the token-ring model is $12,900.

Cayman also bolstered its routing line by adding the first WAN expansion modules to its GatorRoute iR router and GatorStar iHR integrated hub/router. The new RS-232 and V.35 serial line interface modules, which are daughtercards that slide into front-panel expansion slots, complement the existing LocalTalk module.

The Serial V.35 Expansion Card supports dedicated digital connections of up to 115.2K bit/sec, while the Serial RS-232 Expansion Card supports a link to an external modem.

Available this summer, the expansion cards cost $395 each.

©Cisco: (800) 553-6387; Sumitomo: (408) 737-8517; Cayman: (617) 932-1100.

Yes! Unfold the whole story.

Please send me your Lattis System 800™ brochure. And cover everything there is to know about your new eight-port hub.

Name _____

Title _____

Company / Organization _____

Address _____

City/State/Zip _____ Telephone _____

Fax Number _____

≣≣SynOptics

Just because you're looking for a low-end hub doesn't mean you have to look to a low-end vendor.

Introducing the Lattis System 800™ from SynOptics.® It's an eight-port hub, but it still has everything you'd expect from the leader in intelligent hubs: superior quality, a competitive price, and a no-catch, lifetime warranty.

And it's not like those other so-called lifetime warranties. If you read the fine print, you'll see that often the key components like the power supply or fan assemblies are only covered for a limited time. If they're even covered at all.

SynOptics, on the other hand, guarantees the entire unit for life—right down to the nuts and bolts. If anything goes wrong in the first year, we'll replace it. After that, if we can't repair it, we'll replace it. No maybes. No asterisks.

The LattisSphere System.

Where hardware and software work together for a higher level of control.

It's the kind of quality and service you'd expect from a company whose hubs connect more than 10 million nodes.

The Lattis System 800 has plug-and-play connectivity with at-a-glance monitoring for stand-alone networks. And at $399, it has one of the lowest suggested retail prices on the market. The same goes for the 810M model. For only $799 list price, you get SNMP management over IP and IPX network protocols—with support from our Lattis EZ-View™ and Optivity™ network management software.

So don't settle for anything less than SynOptics. Call your authorized SynOptics reseller or 1-800-PRO-NTWK, ext. 204 to learn more. But do it today. Your business warrants it.

1-800-PRO-NTWK
EXTENSION 204.

≡SynOptics
Building the Network Fabric

Other workgroup hubs could cost you a wad.

Hewlett-Packard's lifetime warranty excludes power supply and fan assemblies.
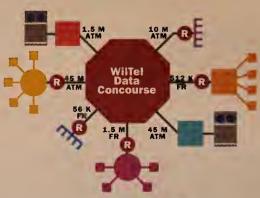
3Com's lifetime warranty excludes power supply and fan assemblies.

Asanté's warranty doesn't cover the life of the product.

WilTel Is Dedicated To Business. Focused On Data. And Offers The Latest Technology With Customized Solutions.

It's plain to see we're the data specialists. For over a decade, WilTel® has introduced innovation after innovation: The first Fractional T-1 service. The first proactive ESF DS-1 performance monitoring. The first frame relay service.

We're also a leader in ATM-based services. And WilTel's Data Concourse will even provide flexible interworking between frame relay and ATM sites.

So call WilTel, your one-stop-shop for business communications solutions, both domestically and internationally. You'll find everything you need for data, voice and video. From a nationwide fiber-optic network to premises equipment to customized LAN-to-WAN integration and operation. As well as a commitment to personalized service. See how easy it is with WilTel.

WilTel's Data Concourse™ is an advanced network platform that translates between different interface protocols. It offers a full range of connectivity from 64 Kbps up to 45 Mbps. Plus interworking between your ATM and frame relay sites. That means you choose the right solution to support your business applications – site by site.

WILTEL
TURNS UP TECHNOLOGY

1-800-364-5113 • Tulsa, Oklahoma
info@wiltel.com
© 1994 WilTel

ARE YOU TRUSTING YOUR DATA TO A VENDOR THAT CAN'T SEE PAST LONG DISTANCE?

by Scott Bradner

The global character of the Internet

Last time, I talked a bit about the scale of the Internet and noted in passing that there are now 78 countries with direct Internet connections. I'm now off to look firsthand at this international aspect of the growing Internet — by the time you read this, I'll be at the Internet Society (ISOC) annual conference in Prague in the Czech Republic.

I'll be joining around 1,000 other network people from, according to the preregistration figures, more than 70 countries; there may be more than 90 represented by the time the meeting starts.

More than three-quarters of the connected countries will be represented, and there will be people from a number of soon-to-be-connected countries. Less than one-quarter of the attendees will be from the U.S.

The ISOC holds a workshop the week before its annual meetings, selecting students from a large pool of applicants. This workshop is designed to provide a foundation of basic networking knowledge and experience to those who will be building the national data networks in their own countries.

The students are then expected to help with the evolution of their local and national data networking.

In many countries, data networking is a fresh idea. Peter Ford, in a presentation at a recent forum put on by the Harvard Kennedy School Center for Science and International Affairs, described four stages in the evolution of national data networking.

In the first stage, the major interorganizational data networks are dedicated to specific scientific endeavors. In the U.S., this stage was pre-National Science Foundation Network (NSFNET) when the Advanced Research Projects Agency Network and the High Energy Physics Network were the prime examples of our national data networks.

Next, the restrictions on the use of the data networks are relaxed and some become general research and education networks. The growth of the NSFNET characterized this stage in the U.S.

The third stage sees the growth of commercial data networks. We are well into this stage now in the U.S., with many companies offering Internet connectivity as part, or all of, their normal course of business.

The final stage, according to Ford, is the common availability of a public Internet. At this stage, internetworking is as pervasive as the phone system.

In most of the world, data networking is in Stage I or II. It is also complicated by the fact that in most countries, the phone system is run by a single public telephone company or PTT. These types of organizations are notoriously slow to understand and adopt new technology. It is hoped that the ISOC workshop will give them a technology injection and help move things along the evolutionary track.

Shifting gears now, there is one more bit of fallout from the actions of those self-anointed defenders of the right to abuse the Internet by posting advertisements to Usenet newsgroups without regard to the topic of the newsgroup. In Prague, the ISOC board will start discussions on a code-of-conduct statement for Internet operators and users. Not that any such document would stop the likes of those now-notorious lawyers from their self-appointed rounds.

I do know a number of good lawyers, but the actions of these two do keep bringing to mind bad jokes.

After my last column on this topic, Don Esry sent me a few hundred lawyer jokes. Here is one of them.
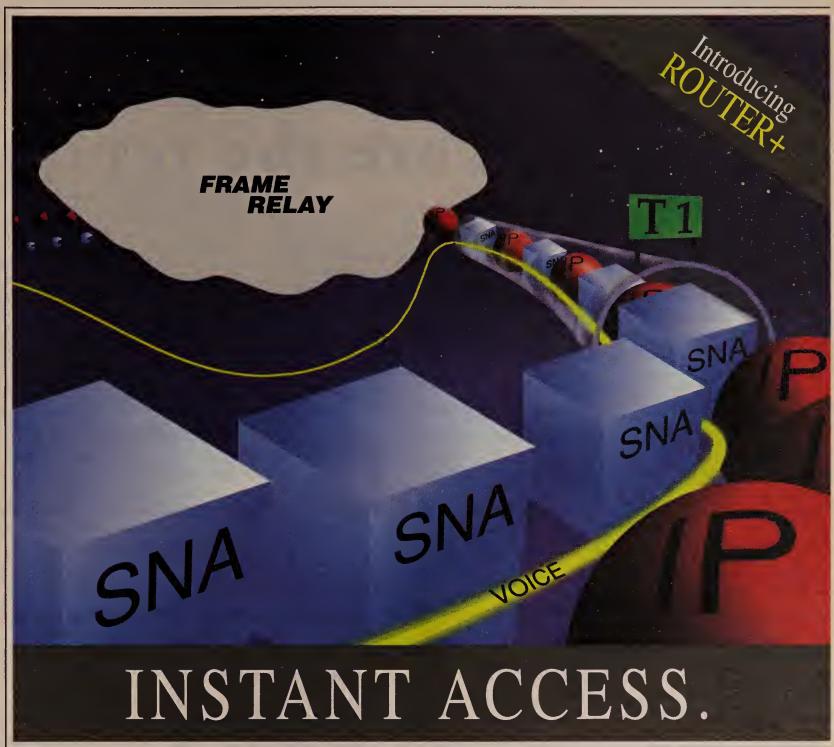
"Why did you switch from rats to lawyers for your biology experiments?"

"We found that lawyers are far more plentiful, the lab assistants don't get so attached to them, and there are some things even a rat won't do."

Disclaimer: I'm an ISOC board member and will be teaching at the workshop; Harvard has no connection with either activity.

➡◆ *Bradner is a consultant with Harvard University's Office of Information Technology. He can be reached via the Internet at sob@harvard.edu.*

P&W set to ship real-time backup for Windows NT nets

BY CHRISTINE BURNS

Chalfont, Pa.

Start-up P&W Technologies, Inc. this week will introduce a utility that provides real-time backup for Microsoft Corp. Windows NT and Windows NT Advanced Server (NTAS) machines.

Octopus is software that sits on a personal computer or server and uses a process called electronic vaulting to back up data across one or more local-area networks.

Electronic vaulting, traditionally a mainframe function, captures file changes as they occur and automatically mirrors the altered data to a local or remote server that is accessible to network users at any time.

According to Robert Gottfried, president of P&W, electronic vaulting makes it easier for network administrators to recover from system failures without any data losses.

"We let users pick which files they want to back up, which means they are only protecting files they need and aren't wasting time and money storing and recovering files they don't," he said.

SOME OVERLAP

Octopus could be used as a stand-alone backup solution but is better suited to supplement existing enterprise tape backup schemes.

"P&W is definitely selling into a market that overlaps with the tape backup vendors," said Richard Buchanan, senior analyst of computing strategy services at Forrester Research, Inc. in Cambridge, Mass.

But Forrester's research has shown that services such as automated backup are going to run on dedicated servers in future client/server systems, he said.

"One of the services that has traditionally been provided by hosts is automated backup. But we believe that in client/server nets there will be classes of servers that are dedicated to things like LAN backup," Buchanan said.

Getting your arms around Octopus

Windows NTAS-based server copy
with a 100-client license $3,975

Windows NTAS-based server copy
with unlimited client licenses $4,975

Windows NT-based desktop copy $395

All versions are now available.

"These dedicated servers will exist at the corporate level and have attached tape libraries to which the company as a whole will back up data stored on individual LANs on a scheduled basis," he added.

The value added by Octopus is that the software is focused on real-time LAN backups "that are easy to get at, which is not always the case with the larger [backup] servers," he added.

Using the Windows user interface shipped with Octopus, a network administrator can select one or more secure locations for real-time file mirroring. This user-defined Guarded File Repository (GFR) can be any storage device attached to a PC or workstation running Windows NT or a server running Windows NTAS.

According to William Whitman, vice president of technology at P&W, the GFR can be either LAN- or wide-area network-based, and can be connected via Ethernet, token ring or X.25.

Octopus requires that Network Basic I/O System be running on both the primary as well as the mirrored servers and workstations, he added.

OCTOPUS OPTIONS

Octopus is a one-to-one, one-to-many or many-to-one data protection system, Whitman said.

In the one-to-many mode, Octopus users can send a single file to multiple locations, providing as many levels of fault tolerance as needed.

In the many-to-one mode, files from disparate client machines are sent to a single place allowing multiple systems to access a central data repository.

In the event the link between two Octopus machines is interrupted, the software notifies the system administrator and begins to log all unsent file activity.

Once the link has been restored, the unsent file activity is then sent to the backup location. If the local Octopus workstation fails, users can either point their applications to the mirrored location and continue as if the disruption had not occurred or wait for the local system to become available and copy the mirrored files back from the remote site.

Gottfried said the next version of Octopus will incorporate an easier way to retrieve data from the mirrored machine.

"Right now, we can tell you where the file is stored, but you have to manually key in the commands to actually restore it back on your local machine," Gottfried said.

He added that Octopus would support other server and client platforms in the future but declined to say which ones.

©P&W: (215) 822-8075.

Optus rolls out first in enterprise fax server line

BY CARYN GILLOOLY

Somerset, N.J.

Optus Software, Inc. this week is expected to bring out the first in a line of enterprise facsimile servers powerful enough to provide faxing capabilities to end users on a variety of LANs.

FACSys Fax Messaging Gateway 4.0 will be based on 32-bit operating systems, such as Microsoft Corp. Windows NT and IBM OS/2. The previous version, FACSys 3.4, only runs on DOS systems.

"We're breaking into the world of 32-bit, multitasking, preemptive operating systems," said Joseph Avellino, president of Optus, based here. "This will let us do quite a bit more than our existing DOS-based products."

The new product will provide faxing capabilities to end users on Novell, Inc. NetWare, Microsoft Windows NT and LAN Manager, IBM LAN Server and Unix-based local-area networks such as Banyan Systems, Inc. VINES. FACSys 3.4 only supports NetWare.

All FACSys software must run on a dedicated machine acting as an application server on the LAN.

The first iteration of the new line, available in September, will run on Windows NT. The software will be a native Windows NT application that will use Optus' own version of Named Pipes to let NetWare clients access fax services. In addition, it will support Transmission Control Protocol/Internet Protocol.

In addition, Version 4.0 includes bindery import facilities to let NetWare administrators manage the servers through the NetWare bindery.

Although Avellino would not specify a time frame, he said Optus will also release OS/2 and NetWare Loadable Module versions.

Enterprise faxing

Projected 1994 fax
server market shares*

Cheyenne — 15%
RightFAX — 14%
Alcom — 16%
Optus — 19%
Other — 23%
Intel — 9%
OAZ — 4%

Total: 2,330 systems

*Includes fax servers with four or more fax ports.

SOURCE: DAVIDSON CONSULTING, BURBANK, CALIF.

GOOD NEWS, SO FAR

For those customers looking to implement Windows NT, the new FACSys software could be a welcome addition.

"We're currently running their DOS-based product, mainly for faxing out reports that are created on the mainframe," said Wade Jenkins, a systems engineer on contract at Northern Natural Gas, a division of Enron Corp. in Houston that has asked to be a beta site for the new Windows NT-based version.

Northern Natural Gas creates reports on a mainframe and sends them to the company's NetWare LAN, where the reports are then faxed out in batch mode. Each batch consists of about 1,600 pages of

---

BRIEFS

Intel Corp. this week will announce a new version of its Storage Express XLC **backup server** for Novell, Inc. NetWare local-area networks. New features include an autochanger for automating traditional tape rotation procedures, tightened integration with Novell's NetWare Management System (NMS) and support for Fiber Distributed Data Interface. Available in June, the StorageExpress XLC costs $11,995.

Upgrade prices for the autochanger, NMS link and FDDI support are $6,495, $2,995 and $2,295, respectively.

Intel: (800) 538-3377.

Attachmate Corp. of Bellevue, Wash., last week announced its new **ZIP Systems Network Architecture** server that gives personal computer and Apple Computer,

Inc. Macintosh users local-area network-to-3270 connectivity. ZIP SNA extends Attachmate's existing 3270 gateway to support Transmission Control Protocol/Internet Protocol and Banyan Systems, Inc. VINES users. Previously it only supported NETBIOS and Internet Packet Exchange networks.

A new ZIP management console sits on top of a ZIP SNA server and adds a Windows-based network management capability for monitoring and controlling gateways.

The ZIP SNA server costs $995, while the ZIP Console costs $1,995. Both products are available now.

Attachmate: (206) 644-4010.

Pittsburg-based **Fore Systems, Inc.** has started shipping its first 155M bit/sec Syn-

chronous Optical Network (SONET)-based **Asynchronous Transfer Mode (ATM) adapter cards** for a variety of personal computer and workstation platforms. Cards are for SBus, VMEbus, Extended Industry Standard Architecture (EISA) and Hewlett-Packard Co.'s version of EISA.

The cards are based on an imbedded Intel Corp. i960 Reduced Instruction Set Computing processor with support for direct memory access and ATM Adaption Layers 3, 4 and 5.

The SBus ($1,895), EISA ($1,995), HP EISA ($2,295) and VMEbus ($2,995) cards are available now and ship with multimode fiber connections. SONET-based cards for Micro Channel Architecture- and NuBus-based platforms will be available by year end, as will cards that support 155M bit/sec over Category 5 unshielded twisted-pair wiring.

Fore Systems: (412) 772-6600.

Retix enters E-net switch fray with stackable offering

Backplane provides new wrinkle in crowded field.

BY SKIP MacASKILL

Santa Monica, Calif.

Retix last week unveiled the first fruits of its recent Calios, Inc. acquisition with the rollout of a new stackable Ethernet switch based on Calios' technology.

The new SwitchStak 5000 is a Reduced Instruction Set Computing-based Ethernet system comprising fixed-port switching modules linked via a high-speed external backplane. Each SwitchStak module supports eight ports of 10M bit/sec dedicated Ethernet via RJ-45 connectors or an attachment unit interface port.

Each eight-port module can support thousands of media access control (MAC) addresses, positioning the modules as lower priced alternatives to routers for linking hubs.

Workgroup Ethernet switches, such as those offered by Grand Junction Networks, Inc., support a single MAC address at each port and are designed to provide dedicated Ethernet links to high-performance workstations.

"Viewing the switch in a backbone manner for segmenting Ethernet LANs is a sensible tactic for Retix to take since it's coming from an internetworking background," said Mary Petrosky, an analyst at The Burton Group, a consulting firm in Salt Lake City. "Its switch family is also compatible with its existing line of routers, allowing users to deploy either technology as needed."

The modules also support two expansion slots, one of which is dedicated to a slide-in StakBus daughtercard. StakBus is a 175M bit/sec external backplane that allows users to stack as many as eight SwitchStaks together and manage the 64-port maximum configuration as a single logical repeater.

The StakBus cards are available in unshielded twisted-pair and fiber versions, as well as a model that supports both fiber and copper connections. This gives users the flexibility to interconnect switches locally or to distribute them as far as 2 kilometers from each other.

The price of switching		
Product	**Price**	**Available**
SwitchStak 5000	$5,400	Now
Fiber-based StakBus	$3,495	Now
Copper-based StakBus	$1,795	Now
Redundant power supply	$495	Now
Fast Ethernet backbone interface	TBA	4Q
FDDI backbone interface	TBA	1st half of '95
TBA = To be announced		

While the initial StakBus is only providing 175M bit/sec of backplane capacity, Retix will develop higher speed StakBuses for better performance and throughput, according to industry observers.

"Just about every networking vendor is embracing Ethernet switching and providing stackable solutions," Petrosky said. "Retix, however, will allow users to upgrade to a higher speed backplane in the future — a degree of flexibility you're not going to get with other stackables."

The second expansion slot supports high-speed interfaces to backbones or servers. The initial expansion modules will support Fiber Distributed Data Interface and 100M bit/sec Ethernet (see graphic, this page).

VIRTUALIZATION

Retix provides a virtual networking capability with SwitchStak, allowing net managers to assign any port in the stack to any one of up to 256 virtual networks. Those virtual workgroups can then be linked together via a backbone router.

This technology enables net managers to create firewalls between different groups of users, as traffic from any port not included in a virtual segment cannot access that segment.

With the SwitchStak rollout, Retix becomes the most recent vendor to use the acquisition route to enter the rapidly growing switch market. Cisco Systems, Inc. (Crescendo Communications, Inc.), 3Com Corp. (Synernetics, Inc.) and Chipcom Corp. (Artel Communications Corp.) all made similar moves over the last several months.

These larger players have a distinct advantage over the few remaining independent switch vendors, according to Petrosky.

"The high-production vendors will have a big advantage over the smaller players because they can pump out product in higher volumes, which will let them drop pricing more quickly," she said.

©Retix: (310) 828-3400.

Imagine A Simple New Way For All Your PC LANs To Play Together.

Now you can direct all your PC LANs simply and easily without leaving your seat. **PATHWORKS**® gives you the first integrated management for multi-vendor PC LANs – NetWare,® LAN Manager, and **PATHWORKS** – from a single PC. Digital's new **PATHWORKS ManageWORKS** is a Windows™ application, totally independent from your network operating system and server. It gives you a single consistent view of all your PC LANs. And it's just one of the many ways **PATHWORKS** helps you integrate the diverse PC LANs proliferating throughout your business. No one knows more about multi-vendor connectivity than Digital. And nothing integrates more PCs, servers, network operating systems, and protocols right out of the box than **PATHWORKS** – the foundation for your business-critical client/server solutions. And now with **PATHWORKS**, you can make them all play together.

CALL 1-800 DIGITAL

Putting Technology To Work **d i g i t a l**

Locus upgrades PC-Interface Unix connectivity software

BY CHRISTINE BURNS

Inglewood, Calif.

Locus Computing Corp. next week will roll out a new version of its PC-Interface personal computer-to-Unix host connectivity software for accessing services across a network.

With new support for a number of enterprise directory services, PC-Interface 5.0 lets PC and Apple Computer, Inc. Macintosh users access services on multiple Unix hosts via a connection to a single Unix system.

According to Dennis Bordelon, a Locus product marketing manager, Version 5.0 of the client and server software includes a gateway that lets users tap into Network File System (NFS), Andrew File System (AFS) and Distributed File System (DFS) networks via a single Unix server.

PC-Interface gets a face-lift

Version 5.0 supports:

▶ A gateway to NFS, AFS and DFS.

▶ Hierarchical host name mapping and domain name service for addressing PCs on Unix LANs.

▶ Simultaneous logons to NetWare and Unix hosts.

▶ IP tunneling for connectivity to remote NetWare servers over TCP/IP.

▶ An SNMP agent.

▶ A new Windows install and configuration utility.

The client component of PC-Interface enables users to immediately access NFS files already mounted on the gateway server, while Locus software at the server gives the clients access to any other NFS servers in the enterprise.

Users gain access to Unix hosts supporting AFS and DFS using the Kerberos authentication scheme. This authentication process occurs transparently to the user when logging on to the gateway for access to AFS or DFS hosts.

"For the client, it means that they don't need to know where they're going [in order] to get there; for the network manager, it means a lot less administrative headaches," Bordelon said.

This single point of access to multiple Unix servers is exactly what users are looking for, said Bill Campbell, president of Celestial Software, Inc., a systems integration firm in Seattle. Campbell is running beta copies of PC-Interface 5.0 at Celestial and has installed it at several customer sites.

"This type of connectivity is a lot easier to use than something like [SunSelect, Inc.'s] PC-NFS because I can mount a Unix system as a drive, and then I have access with that one mount to anything that is available on that system" or on the enterprise, he said.

MORE ADDITIONS

In addition to this enterprise gateway support, PC-Interface 5.0 gives users a way to send and receive electronic mail from any Unix machine in the enterprise as well as the ability to exchange Internet mail.

Administrative features have also been included in Version 5.0, making it easier for net managers to install the software and configure PC-Interface network nodes. The features include a Windows-based configuration tool and support for hierarchical host name mapping for organizing network nodes into logical, physical and arbitrary groups.

PC-Interface 5.0 also supports the Unix-based Domain Name Service (DNS), which lets Windows users connect to named Unix servers using the domain name convention scheme.

Bordelon said that PC-Interface 5.0 is more tightly integrated with Novell, Inc.'s Transmission Control Protocol/Internet Protocol stack than previous versions of the software, which makes it easier for PC users residing on NetWare networks to access Unix hosts via a single logon.

Other than the standard new features shipped with PC-Interface 5.0, Locus also offers an add-on functionality package in PC-Interface Plus 2.0. These additions include Eudora, a Windows-based Unix E-mail interface from Qualcomm Software, Inc., and TinyTerm, an advanced terminal emulation program from Century Software, Inc. that provides emulation services for about a dozen terminal types.

PC-Interface 5.0 and PC-Interface Plus 2.0 are available for $279 and $49, respectively.

©Locus: (310) 670-6500.

Comments?

See "Contacts" box on page 2.

Novell launches NetWare Ready product integration program

Cabletron, 3Com among vendors to sign on to plan.

BY CARYN GILLOOLY

San Jose, Calif.

Novell, Inc. last week expanded beyond its "Yes" program for testing and labeling NetWare-compatible products by establishing a new program called NetWare Ready.

The NetWare Ready label will be given to those products that come bundled with NetWare or NetWare-related products.

This goes beyond the Yes program, which simply ensured that other vendors' products had been tested to work with NetWare.

Toby Corey, senior director of marketing for Novell's NetWare Products Division here, said an example of a NetWare Ready product would be a Zenith Data Systems Corp. laptop computer that ships with the NetWare client and includes prebuilt configuration files.

"[If I am a customer,] I don't need to call the MIS guy to get the client piece, and I don't need to install the right driver," Corey said. "It's already there."

CRITICAL SUPPORT

Third-party vendors are already rallying around the NetWare Ready program.

Such heavy hitting companies as Cabletron Systems, Inc., Compaq Computer Corp., Digital Equipment Corp., SynOptics Communications, Inc. and 3Com Corp. last week announced that they were all working with Novell to bundle pieces of NetWare in their respective products.

Cabletron, for example, announced that it will begin offering NetWare and related Novell products preinstalled within its PC Media Interface Module (PCMIM) for its Multi Media Access Center wiring hubs.

PCMIM is essentially a personal computer within a hub. It is an Intel Corp. 486DX/2-based processor that lets customers load applications — such as management, routing and communications software — onto the hub rather than in on a separate PC attached to the hub.

> "[If I am a customer,] I don't need to call the MIS guy to get the client piece, and I don't need to install the right driver," Corey said. "It's already there."

Initially, PCMIM is available preinstalled with NetWare 3.X or NetWare 4.X as well as with NetWare NFS, NetWare for SAA, NetWare HostPrint, NetWare SNA Links or NetWare Management System (NMS) applications.

3Com, which recently announced a product called AutoLink for letting customers automatically install NetWare client software, last week said it will bundle the software into its EtherLink III Parallel Tasking and TokenLink III adapters.

In addition, Compex, Inc., which sells FreedomLine network interface cards, last week said it will include Novell's Universal NetWare Client free with its hardware. Even with the software preinstalled, the cards will continue to start at $59 per card.

©Novell: (800) 453-1267.

Optus

landscape faxes and takes as many as eight hours to complete.

"We can get much higher output with the NT version — there's a lot of coprocessing that can happen that we can't get now," he said.

"We'd like to double or triple our [faxing] speed, and based on what we understand about the product, we think we'll be able to do that," he added.

Analysts, however, were skeptical about whether initially supporting Windows NT rather than a more proven 32-bit operating system is the best strategy for Optus.

"There's absolutely a need for more powerful, multitasking fax servers — the idea of servicing a large department or even an enterprise on a DOS-based machine is really inefficient," said Pete Davidson, president of Davidson Consulting, a fax consulting and publishing company in Burbank, Calif.

"But whether NT is the right platform is the $64,000 question," he added.

Avellino defended the company's position: "We think NT will be the winner [over OS/2 and Unix] as *the* application server."

Pricing is not yet available.

©Optus: (908) 271-9568.

NET RESULTS

by Mark Gibbs

When network computing is really not work computing

Do you remember the heady days of the late '80s when all of the network operating systems vendors, notably Novell, Inc., went around talking about network computing?

Today, the term lingers on and is used not only as if the capability was possible, but even desirable.

Network computing once seemed to be an attainable goal. The idea was to get all of an organization's computing resources networked into one great, oiled, smoothly running machine.

Ha! We're no nearer to that goal about a decade later.

Worse still, no one ever said exactly what network computing was supposed to achieve. It was presented as being a gestalt — the sum of the parts was greater than the whole.

But look at any level of the net from the desktop to the server, and the problems of implementing net computing are apparent.

For a start, we have an unspeakable mess on the desktop. We have a collection of computing architectures there that are, at best, user-centric and, at worst, based on somebody's good idea.

There are few user-centric systems; the Macintosh is the only major one that leaps to mind.

The good idea systems — notably DOS and Unix — had networking added as a system service. But that service turned out to be as user-unfriendly as the rest of the system.

To get the resources on the desktop to work as part of the network rather than what we have — the network resources supporting the desktop — pretty much requires throwing out the desktop systems we have. And that's a wholly unthinkable idea.

At the server end, we have another mess. It doesn't matter whether you go with NetWare, Windows NT Advanced Server or VINES because you still get a server-centric view of the network.

"Howsabout Macintosh System 7, LANtastic or Windows for Workgroups?" you might ask. Sorry, but they just do the same thing at a finer granularity.

Even with the global directory services of VINES and NetWare 4.X, you only gain a sophisticated inventory of server resources and network objects, rather than a truly integrating principle.

Now let us consider what is in the middle. We have bridges, routers, print servers, network-compatible printers, hubs, communications servers and a score of other devices. But not one of these device types has a standard view of the world; they each support different and specific network environments.

So what's my point? Simply that network computing will never happen because it is not practical or possible, and because no one needs or wants it.

It isn't possible because the same diversity of products that prevents computing resources from working together is what keeps the network industry rushing forward.

But the main reason net computing will never happen is that no one needs or wants it. That's because net computing buys you nothing.

It doesn't make better use of resources given that they are already cost-effective. Any architecture that would try to improve on existing network resources would have greater overheads than benefits.

No one wants network computing because networking is a service and not an end in or of itself. Networking is about supporting personal productivity.

In our race for networking sophistication, we often forget that what matters most is what happens on the desktop. If the user isn't supported and serviced adequately, the network is failing.

We need to expunge the concept of network computing from our vocabulary.

Another piece of marketing gobbledygook bites the dust.

Gibbs is a consultant and writer in Ventura, Calif. He can be reached at (800) 622-1108, Ext. 504, or on the Internet at mgibbs@rain.org.

How Did World Airways Re-route Their PC Network To TCP/IP?

The answer: PC-NFS® from SunSelect, the world's most popular software for connecting PCs to TCP/IP networks. Whether you use MS-DOS® or Microsoft® Windows, you can connect to any UNIX,® IBM® or VMS™ environment for a true open system solution. And because SunSelect is a Sun Microsystems, Inc. business, you know you're getting the best with PC-NFS.

An air traffic controller's worst nightmare.

World Airways' old mainframe "hub" was struggling to handle the company's large volume of PC user traffic: an ad hoc query typically took *over three days*. Unhappy users were stuck in a holding pattern, and productivity was grounded.

World Airways needed to replace the old mainframe and upgrade to a UNIX client/server solution that could link every computer directly to the servers. And that solution included PC-NFS software: It connected all 150 PCs and 26 remote users to a new TCP/IP network in a snap.

PC-NFS
The TCP/IP network solution for PCs.

Up and flying in under a week.

PC-NFS saved World Airways time, because it was fast and easy to install; money, because it utilized existing hardware; and overhead, because it reduced system costs by 45%. More importantly, PC-NFS dramatically *increased* PC user productivity by providing PC desktop access to a wide variety of UNIX-based payroll, accounting and air transport applications – for example, PC-NFS let employees use the UNIX network as a gateway to on-line services like Lexis™ and Nexis™ And because it's completely transparent, users can't tell whether they're working locally or on the server.

PC-NFS opens a world of opportunities to your PC users, too. Through seamless TCP/IP connectivity, PC-NFS users can access enterprise resources – like applications, network printers, e-mail and more – without sacrificing their familiar Microsoft Windows or DOS interface. PC-NFS offers:
• Full integration with DOS, Microsoft Windows, NFS and any TCP/IP network;
• Powerful, full-featured applications like FTP and Telnet,® plus on-line help;
• Broad support for Windows Sockets and NetBIOS applications, Microsoft Windows for Workgroups,™ SelectMAIL,™ RUMBA® for PC-NFS and more.

FREE facts on connecting PCs to TCP/IP.

Call **1-800-60-SELECT, ext. 0304** today to get our FREE booklet, *"10 Things You Must Know Before Connecting PCs to a TCP/IP Network."* It describes important considerations for successful right-sizing to a TCP/IP network. So call now and start re-routing your network to TCP/IP with PC-NFS.

❋ SunSelect
A Sun Microsystems, Inc. Business
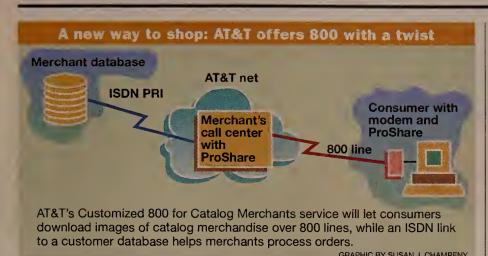
Call 1-800-60-SELECT for a FREE copy of "10 Things You Must Know Before Connecting PCs to a TCP/IP Network."
EXTENSION 0304

GLOBAL SERVICES

A new way to shop: AT&T offers 800 with a twist

Merchant database

ISDN PRI

AT&T net

Merchant's call center with ProShare

800 line

Consumer with modem and ProShare

AT&T's Customized 800 for Catalog Merchants service will let consumers download images of catalog merchandise over 800 lines, while an ISDN link to a customer database helps merchants process orders.

GRAPHIC BY SUSAN J. CHAMPENY

AT&T poised to deliver on-line 800 catalog service

BY JOANIE WEXLER

The first in a series of AT&T 800-service packages customized for vertical markets will be announced later this month in the form of an on-line catalog bulletin board service, according to sources familiar with the carrier's plans.

AT&T has said it will differentiate itself in the 800 market by putting together packages of services honed for eight types of businesses: retail, general merchandising, banking, insurance, health care, travel, chemical and trucking (*NW,* May 2, page 7).

One package, expected to be available this summer, is Customized 800 for Catalog Merchants. AT&T will reportedly allow merchants to use the AT&T net as a bulletin board from which consumers can download and view images of catalog merchandise and place orders.

The concept is very much like the on-line ordering capabilities available today via America Online, Inc. and Prodigy Services Corp. bulletin boards, analysts said.

Merchants electing to use the AT&T service to market their goods would outfit their call centers with a modem and Intel Corp. ProShare Standard Edition software. The ProShare package allows personal computer users to exchange voice and data, including color pictures, with other terminals similarly equipped, according to Intel.

This means that consumers dialing in to the merchant's catalog would require a modem and ProShare Standard Edition.

ProShare Standard Edition lists for $99, although Intel said at the ComNet '94 conference in January that companies purchasing ProShare and AT&T service together would be eligible for discounts.

Merchants would also run an Integrated Services Digital Network Primary Rate Interface link from their 800 call centers to their corporate databases. ISDN's automatic number identification feature would let computers in the merchant net quickly access customer profile and ordering history information from the database.

Sprint talks with int'l carriers

France Telecom and DBP Telekom looking to buy into Sprint.

BY DAVID ROHDE

Kansas City, Mo.

Finally, a deal that's likely to go through.

Sprint Corp. last week confirmed that it is negotiating with France Telecom and Deutsche Bundespost (DBP) Telekom for the two European carriers to jointly purchase a 15% stake in Sprint.

In a statement, Sprint said "the parties have reached concurrence on many principles, [although] some issues remain."

By contrast, the recently failed discussions of a proposed "merger of equals" between Sprint and EDS Corp. involved a complicated stock swap and hinged on EDS being divested by its parent, General Motors Corp.

Based on the kind of prices BT paid for a stake in MCI Communications Corp., Sprint would garner $4 billion in the transaction, said Jeanine Morley, a research analyst with New York brokerage firm Oppenheimer & Co.

"They could go ahead and build the type of information systems they were looking at with EDS based on the $4 billion they're getting," Morley said.

For example, Sprint could use such a cash infusion to speed up deployment of its Synchronous Optical Network backbone using Asynchronous Transfer Mode switching. That means Sprint customers could see benefits out of the deal domestically before any changes come about internationally.

"The more they invest, the more likely they are to make progress," said Ernie Smith, senior vice president of information systems for Sprint customer Advest, Inc., a brokerage firm in Hartford, Conn. "Besides, I like the idea of them becoming a bigger global player."

MCI, with greater ratios of debt than Sprint, had to spend its BT proceeds on debt repayment.

Sprint could even use the proceeds to reopen discussions with GM about purchasing EDS outright, rather than a merger of equals, said Berge Ayvazian, senior vice president of The Yankee Group, a Boston-based consultancy.

"Now Sprint will gain the upper hand," he said. "I think it would be very hard for GM to walk away, given the new factor in this deal." Sprint's discussions with

France Telecom and DBP Telekom probably preceded its discussions with EDS, he added.

Another possibility for Sprint is to use the additional capital to invest in personal communications services licenses, scheduled for auction by the Federal Communications Commission, according to an Oppenheimer report.

Overseas, the deal would help Sprint solve a visibility problem in foreign markets, analysts said.

"It will help Sprint catch up in promoting its local identity and local presence in the French and German markets," said Caroline Michel, an analyst at International Data Corp., a research firm in Framingham, Mass. Until now, Sprint has sold its switching platforms to different operators but has failed to gain wide exposure with corporate clients in those two countries, she added.

France Telecom and DBP Telekom last December formed an alliance — Eunetcom, counterbalancing another European alliance — Unisource (see graphic). Unisource, in turn, reportedly has been discussing an alliance with AT&T.

"AT&T had been in discussions with France Telecom and DBP Telekom last year, but they were unable to come to terms with [them]," Ayvazian said.

Added Peter Bernstein, research vice president with consultancy Probe Research, Inc. in Cedar Knolls, N.J., "BT had threatened to go it alone globally and plop down some switches [in foreign countries], and AT&T has made some noises about doing the same thing." ∎

Jockeying for position

U.S. carriers are teaming up with the following international carrier alliances:

Unisource
- PTT Netherlands
- Swiss PTT
- Telia International (Sweden)

Eunetcom
- Deutsche Bundespost Telekom (Germany)
- France Telecom

WorldPartners Association
- AT&T
- Kokusai Denshin Denwa Co., Ltd. (Japan)
- New Zealand Telecom
- Networks & International Ltd.
- Singapore Telecom
- Telstra OTC, Ltd. (Australia)

Marc Ferranti and Martin LaMonica, IDG News Service correspondents, contributed to this story.

BRIEFS

Worried that security concerns will stall development of the **National Information Infrastructure (NII),** the **Clinton administration** has scheduled a public meeting on NII security in Washington, D.C. on July 15. Those wishing to speak at the meeting must submit a one- to two-page position statement by June 28 to NII Security Issues Forum, c/o U.S. Treasury Department, 3090 Annex, Washington, D.C. 20220. The statement may be faxed to (202) 622-2057 or sent via Internet to nii.security@treas.sprint.com.

MCI Communications Corp. last week said it is making available two prepaid debit cards — telephone calling cards with predetermined amounts of calling time that are paid for in advance. One, **MCI PrePaid Card,** reportedly will allow companies to better manage long-distance calling costs by providing traveling or temporary employees with ceilings on their spending limits and limiting losses associated with lost or stolen cards. The other, **MCI Promotional Card,** is for companies to give to their customers and prospects; a special message can be played each time the user employs the card.

MFS Communications Company, Inc. continued its expansion last month with the announcement of new fiber networks in **Phoenix** and **San Diego.** The two cities will bring

MFS' tally of metropolitan areas served up to 26, including London. The carrier is also expanding its 123 route-mile Chicago network to serve suburban business centers in Naperville and Northbrook.

Associated SMR, Inc. and other shareholders of Corporacion Mobilcom, S.A. de C.V. said last week that they have signed a letter of intent with **specialized mobile radio (SMR)** service provider **Nextel, Inc.,** under which Nextel would invest $165 million in return for a 22% equity interest in Mobilcom. **MCI Communications Corp.** invested heavily in Nextel earlier this year as its portal into the wireless market.

Mobilcom and its affiliate Radiophone, S.A. de C.V. hold interests in corporations with SMR licenses in major metropolitan

areas in **Mexico** — a potential avenue for MCI in that country.

AT&T said it has decreased its prices for **International Accunet T1.5** T-1 services to **Mexico** by up to 19% in the 0 to 50 and 51 to 100 mileage bands.

AT&T Network Systems and **BroadBand Technologies, Inc.** have won a contract to extend fiber to the curb for **Southwestern Bell Telephone Co.** in Richardson, Texas. Southwestern Bell wants to provide 47,000 consumers in Richardson with video on demand and interactive shopping and educational programs — if it wins the right to provide video services. Without such approval, the carrier may be limited to providing telephone services over the broadband infrastructure.

by Eric Paulak

Local contracts are coming...slowly

I f you're the biggest customer in the market, you get the best deals, right? That's the way AT&T, MCI and Sprint play it with their contract tariffs, Tariff 12s and Specialized Customer Arrangements.

But what about the local level?

Traditionally, the only contract deals you could get from the Bell companies were for Centrex and local high-capacity lines such as T-1s and some T-3s. But that's now changing.

NYNEX has its first contract deal for measured service in Massachusetts. While the tariff doesn't give you all the details you can get from an AT&T or MCI contract, it's a start, and other carriers should follow suit.

NYNEX has introduced what it calls a Customer Specific Pricing Plan (CSP), a feature that allows the carrier to offer intra-LATA toll service at negotiated rates in response to a competitive bid in Massachusetts.

What's nice about NYNEX's CSPs is they tell you who is getting the great deal, unlike long-distance contract tariffs and AT&T's Tariff 12. In this case, it was Raytheon Corp. that inked the agreement.

For a three-year commitment of an undisclosed volume, Raytheon is getting Measured Telecommunications Service (MTS) for $.043 per minute during all time periods. The regular tariffed rates are $.095 for the first minute during the day and $.085 per minute thereafter.

So for a 10-minute daytime call, Raytheon pays $.43, whereas any other user pays $.86. That's a 50% savings.

The only way for other users to come close to those rates under the standard tariff is by committing to 1,500 minutes a month. With that, you'll get a $.03 per-minute discount — after the initial minute — for all calls exceeding the 1,500-minute minimum.

But even with the discount, a 10-minute daytime call still costs $.59. So Raytheon is still saving 27%.

What are other states doing with contract deals? Not much, other than for high-capacity lines and special construction arrangements. BellSouth is pumping out a bunch of both of these types of contract deals, and Pacific Bell has more than 500.

But switched, measured services are still lacking.

Why? The various state public utility commissions (PUC) are holding them up so that rural and small businesses don't end up paying the highest rates.

Pacific Bell could change all that. The carrier has a proposal pending with the California PUC that would allow it to offer all business services on a "competitive" basis, meaning contract deals.

The proposal is being delayed, however, as the PUC sorts out the overall issue of local competition and some of the requests Pacific Bell has made in response to pending competition.

One such request is that contract deals be allowed to take effect when they are signed and that they no longer be required to carry the California PUC's approval. Another proposes to eliminate tariff requirements for contract deals altogether. Both of these would be good for the individuals signing the deals, but they hide the rates from the rest of the business users.

Neither proposal has much chance of getting passed — they're just delaying the opening of the local market and the offering of local contract deals.

But once the intra-LATA toll market is open to competition, Pacific Bell will be more than ready to roll out its contract deals. It was ready to sign several dozen contracts last year when California was originally scheduled to offer intra-LATA toll competition. Now it should have even more waiting in line.

And once the deals start coming in California, you can expect more states to follow suit. Until then, take a look at some of the volume and term discounts the Bells are offering. They're not contract deals per se, but they will give you close to the same results.

➥ Paulak is associate publisher for the Center for Communications Management Information in Rockville, Md., a provider of rate and tariff information. He can be reached at (301) 816-8950, Ext. 327.

McCaw certifies wireless gear; industry players form interoperability consortium

BY JOANIE WEXLER

Kirkland, Wash.

McCaw Cellular Communications, Inc. last week took steps to reassure potential Cellular Digital Packet Data (CDPD) users that their portable devices will operate smoothly with McCaw's AirData service, still in pilot stage.

McCaw has successfully tested equipment from several wireless vendors, including modem makers Cincinnati Microwave, Inc., Pacific Communication Sciences, Inc. (PCSI) and Sierra Wireless, Inc., for use on its net.
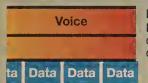
How CDPD works

Carriers have two ways of passing digitally encoded data over the analog cellular net. The CDPD spec accommodates interoperability between both schemes.

Channel hopping: Idle time between voice conversations on different channels is used to pass a single data stream; capacity-efficient, but can cause interference with voice.

Dedicated data channel: No interference, but decreased capacity for delay-sensitive voice.

GRAPHIC BY SUSAN J. CHAMPENY

The move comes on the eve of an announcement that a group of industry players has coalesced to further spur CDPD implementation agreements. Dubbed CDPD Forum, Inc., the group will officially introduce itself at the Wireless DataComm conference later this month in San Jose, Calif.

The forum includes founding carriers AirTouch Cellular Services, Ameritech Cellular Services, Bell Atlantic Mobile Systems, Inc., GTE Personal Communications Services, McCaw, NYNEX Mobile Communications and Southwestern Bell Mobile Systems, industry sources said.

IBM confirmed that it is joining as a charter member, and most wireless equipment and software vendors are also expected to be on the forum's roster.

The group's goal is to "broaden CDPD at all levels, from manufacturers to software providers to carriers," said Russ Brankley, CDPD product manager at Bell Atlantic Mobile and the carrier's representative to the forum. "Interoperability is a large part of that."

Such collaborative efforts as the birth of the Forum and McCaw's interoperability tests are especially needed in the wireless world, which today exists largely as a hodgepodge of devices and emerging technologies that may not mix and match, analysts said.

For instance, "McCaw does not have a presence in all our markets, so...it's very important to have a CDPD unit [user access device] that works in one city work in another and not have to make any changes to it," said Lance Crawford, senior engineer in American Airlines, Inc.'s Sabre Computer Services unit.

American, for example, which has been piloting the AirData service since December, will be served by GTE Cellular in its home office region in Houston, not by McCaw.

American has been testing AirData for communications between wireless terminals at airports throughout the country and American's Tulsa, Okla., data center.

The tests so far have shown "CDPD to be very reliable," Crawford said, although response times have increased from less than a second to between one and five seconds using wireless technology instead of land lines.

The McCaw interoperability efforts technically only show that the devices tested work with AirData, headed for commercial rollout in 105 markets by the end of the year. McCaw's target cities for initial commercial implementation are Dallas; Kirkland, Wash.; Las Vegas; and Miami.

However, BAM, which announced commercial service in April (*NW*, May 2, page 47) has also certified that the Cincinnati Microwave and PCSI modems work with its network, indicating that the devices should be portable across BAM and McCaw networks.

Jeff Brown, vice president of sales and marketing at McCaw's Wireless Data Division, explained that the goal is for a "McCaw user to pick a device up in Washington, fly to Chicago, turn the device on and have it work, even if it's over a different service."

In this example, he said, the Ameritech and McCaw nets will ultimately link "and I'd bill the user for time on the other carrier's network" at McCaw rates.

But Brown said McCaw could not guarantee service quality outside the confines of the CDPD specification in such a scenario. ∎

Start-up gives PacBell the hook over access problem

BY BILL BURCH

San Francisco

A new Internet access provider based here recently found out the hard way not to take local access for granted.

Start-up Hooked, Inc. wanted to offer Bay Area users a sophisticated Internet access service with a graphical user interface, but problems with securing local access from Pacific Bell forced the company to ask competitive access provider (CAP) Teleport Communications Group, Inc. for help.

Although Teleport was able to devise a limited solution for Hooked, regulatory barriers kept the CAP from providing a more permanent fix.

Hooked's story illustrates that the CAPs' entry into new markets is pressuring the local exchange carriers (LEC) to improve service, but regulatory barriers prevent the CAPs from competing with the LECs in more service areas.

To support its access service, Hooked wanted to run two T-1s supporting 48 simultaneous sessions into a U.S. Robotics, Inc. hub that would separate out individual digital data streams.

Pacific Bell proposed Centrex service for the link, which sounded OK to Stan Christensen, Hooked's technical director. It was March, the equipment seemed close to being ready, and Hooked personnel began to attend trade shows to raise the company's visibility.

Then the problems began. Pacific Bell could not get its T-1 lines to work with the U.S. Robotics system. In order for its error correction and other features to work properly, the system needed to use multifrequency signaling, but the carrier used the newer and more common dual-tone multifrequency (DTMF) signaling.

Pacific Bell engineers worked on the problem for a week. "We got a lot of different [explanations] why we couldn't get it," Christensen said. " 'We don't provide it, we can't get it to you, and talk to somebody else.' "

The carrier now explains that its sales personnel initially overlooked the request for multifrequency signaling, noting that DTMF is typically used with end-user equipment.

Unable to get the setup to work, Pacific Bell offered Hooked a credit for the work done to date, then said it could try a different approach, but Hooked would have to foot the bill for it, Christensen said.

Meanwhile, David Holub, Hooked's vice president of operations, phoned around to see if other carriers could devise a solution.

Enter the savior. Teleport had run a fiber ring around San Francisco that came within 50 feet of Hooked's office. More importantly, the carrier had installed an AT&T 5ESS switch and connected it to a Pacific Bell end office in preparation for the state's upcoming launch of intra-local access and transport area toll competition.

> " 'Soon' in PUC time is kind of like the difference between people years and dog years."

That switch allowed Teleport to carry switched data traffic from the Pacific Bell network to Hooked's office. The CAP was able to get a T-1 connection up in four days, and Hooked was on-line in mid-April. From the start, the connection worked flawlessly, Christensen said. Today, the company has roughly 300 users coming in over 24 lines.

The company's future is still clouded by network problems. Teleport is not authorized

to provide switched access in California. To skirt the law, the carrier is not charging Hooked for its T-1, avoiding the legal definition of a "service."

In the meantime, Teleport has filed with the public utility commission (PUC) to offer switched transport to gateway service providers. Although the CAP has asked for expedited treatment, it could take months before permission is granted, according to Michael Morris, Teleport's regional director for regulatory and external affairs.

" 'Soon' in PUC time is kind of like the difference between people years and dog years," Morris said. The state was to have already begun a local competition proceeding in preparation for a service launch in 1997, but nothing's under way yet, he added.

Some of California's other on-line service providers have developed a certain fatalism about local access connections.

The WELL bulletin board service of Sausalito, Calif., uses 75 inbound Centrex lines to serve approximately 9,800 users. When the company had a problem with line noise, fixing it took multiple calls to Pacific Bell, according to Phil Reese, The WELL's customer support manager.

When lines go down at The WELL, the company just waits for repairs, added Reese, who seems resigned to problems in the local loop. "PacBell works hard at what it does, and they still fail — and they're a big company," said Reese, who doubts that smaller companies such as the CAPs could do any better.

Back at Hooked, the threat of Teleport's entry into another of its markets has Pacific Bell engineers scuttling around to get a link in place, according to Christensen.

Even if the carrier can get a link to work, he doesn't know if he can trust the connection and remains torn on whether to return to Pacific Bell.

"We were down for weeks; we couldn't get users on our system," Christensen said, adding that he wonders whether the carrier will reimburse his company for lost business. ⚡

Catalog service
Continued from page 29

"Users would dial up the merchant's bulletin board via the merchant's 800 number and download pictures of merchandise and ordering information," said an analyst familiar with the plans who requested anonymity. Then, if a transaction takes place, a record of it is tied with the customer's new or existing file on the corporate database via the ISDN link, he said.

AT&T will also provide some customer billing services for the merchants, according to the source.

An AT&T spokesman said only that the company has nothing to announce at this time.

A spokesman for JCPenney Company, Inc. said that while the retailer has no current trials in place with the new AT&T service, "that could change next week. We're exploring literally every area of electronic retailing."

One alternative will emerge this summer with the kickoff of AT&T's PersonaLink Services (*NW,* Jan. 10, page 1). The service, which is based on software from Apple Computer, Inc. spin-off General Magic, Inc., will, among other things, provide an electronic shopping mall where merchants can collectively hawk their wares.

Any number of merchants will be able to buy space in the three-dimensional "mall," and users of PersonaLink Services will be able to "go" to the mall and browse from store to store, viewing and purchasing merchandise.

The difference in the PersonaLink electronic mall and the customized 800 service is merely in the type of distribution channel it is, said one analyst. "With the PersonaLink service, merchants share their channel with others."

There is no new technology involved with the 800 for Catalog Merchants service, analysts said. It is merely a way for AT&T to repackage capabilities for a class of customers that would like 800 customization and a way for AT&T to differentiate itself in a competitive 800 market. ⚡

Comments?

IETF shaping multifaceted messaging spec

BY DAVID ROHDE

Arlington, Va.

The Internet Engineering Task Force (IETF) is taking up where a group of vendors left off in trying to forge a standard that supports the exchange of voice, E-mail and fax among dissimilar systems.

The Universal Messaging Interoperability Group (UMIG), formed last year to develop a digital standard to exchange voice, electronic mail and facsimile messages among multiple vendors' systems, disbanded last month.

However, the IETF is pursuing the same goal with a proposed "profile," now out for comment, said Carl Schoeneberger, vice president of Octel Network Services in Dallas. The profile is based on components of the Transmission Control Protocol/Internet Protocol suite such as the Simple Mail Transfer Protocol and the Multi-purpose Internet Mail Extensions (MIME) protocol.

MIME "provides agreed-upon conventions for carrying multiple messages inside of one message," Schoeneberger said.

UMIG was designed to replace the Audio Messaging Interchange Specification (AMIS), a four-year-old analog spec that can transfer only nine voice messages at a time from one system to another and cannot replicate such features as urgent-message flags and spoken mailbox names.

Audio messaging standards history

Sept. 1988	AMIS is formed to develop voice mail interoperability spec.
April 1990	Version 1 of AMIS is released.
Feb. 1992	AMIS Version 2 is released; still supports only low-volume message exchanges.
Sept. 1993	AMIS is disbanded; UMIG is formed to develop a more robust standard integrating voice mail and E-mail.
May 1994	UMIG is disbanded and its work is transferred to EMA.
June 1994	Proposed specs sent out as Internet draft standard.
July 1994	EMA committee expected to start working on UMIG specs.

SOURCE: CAUGHMAN CONSULTANT SERVICES, ATLANTA
GRAPHIC BY SUSAN J. CHAMPENY

UMIG's voice messaging vendor participants blamed the group's dissolution on an inability to interest E-mail vendors and users.

"They were trying to draw the attention of the E-mail community — with not very much success," said Alison Caughman, president of Atlanta-based Caughman Consultant Services, which has facilitated AMIS and UMIG discussions through the Information Industry Association.

"An attempt was made to get LAN organizations to join the group, and those people said, 'We're not all that interested in that right now,' " Schoeneberger said.

UMIG's work has been referred to a committee of the Electronic Messaging Association, based here. It's expected to take up the task next month.

Meanwhile, users are taking another look at AMIS analog.

"People who didn't think they would use analog did," Caughman said. "Now you find people in their RFPs saying that they want to have that capability."

The nine-message restriction — although subject to much criticism — was deliberately designed so that complicated security features would not be required to protect against voice junk mail and hackers.

"AMIS analog is working as designed," said Don Van Doren, president of Vanguard Communications Corp., a voice processing consultancy in Morris Plains, N.J. "The usage is picking up."

A key factor, Van Doren said, is that several private branch exchange manufacturers have incorporated addressing features into AMIS analog that give users within an organization a standard way to dial voice mailboxes in another office or organization.

A separate AMIS standard — AMIS digital — has suffered from the fact that it is based on a 10-year-old version of X.400 and, as a result, has not drawn support from standards-setting bodies. And PBX manufacturers remain skeptical of the prospect of a digital standard emerging from among the major voice mail vendors, even through an Internet effort.

"How many of them are going to open up their channels to do digital?" asked Ken Hu, senior product manager for call centers and voice processing at InteCom, Inc., a PBX manufacturer in Dallas.

AMIS digital today is being used only by three voice messaging vendors — Boston Technology, Inc . Digital Sound Corp. and Unisys Corp. — in sales to regional Bell holding companies for their public network-based voice mail products. ⚡

DON'T GET AN ATTITUDE. IT'S NOT AS IF IT TAKES TWO OF YOU TO PLUG IN A CRAY BOUNDARY ROUTER.

S. Adams

© 1989 United Feature Syndicate, Inc.

BR6100 Ethernet router

BR6500 Ethernet/ token ring router

Cray Communications' new BR6000 "plug and play" boundary routers let you extend network services out to sales people, telecommuters, and even your most remote offices. TCP/IP, IPX, X.25, frame relay, and PPP software is already installed. The user just plugs the router in at the remote site, and you configure it from your central location.
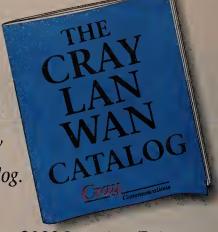
Available in Ethernet and token ring versions, the BR6000 features flash PROM updates and SNMP compliance. Auto dial backup senses a failure or potential failure and automatically establishes a dial-up connection.

A unique dial-in port allows remote users to dial into the boundary router and access the corporate network over the WAN link. This helps reduce long distance charges and gives remote users

full access to network resources.

Cray offers the broadest lines of both LAN and WAN products available from any supplier. All supported by local Cray operations in fifty countries on five continents.

For more information, send for a **FREE** copy of *The Cray LAN-WAN Catalog*. Write to Cray Communications, 9020 Junction Drive, Annapolis Junction, Maryland 20701, or call **1-800-FOR-CRAY** (1-800-367-2729).

THE CRAY LAN WAN CATALOG
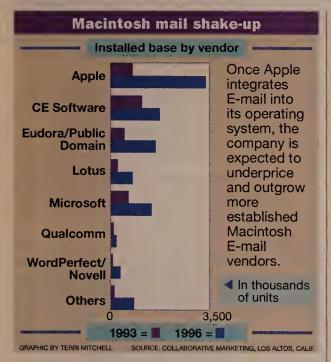
Cray
Communications

BOUNDARY FREE REMOTE ACCESS

CLIENT/SERVER APPLICATIONS

Distributed Databases, Messaging, Groupware, Imaging and Multimedia

Macintosh mail shake-up

Installed base by vendor

Apple
CE Software
Eudora/Public Domain
Lotus
Microsoft
Qualcomm
WordPerfect/ Novell
Others

◀ In thousands of units

0   3,500

1993 = ■   **1996 =** ■

GRAPHIC BY TERRI MITCHELL   SOURCE: COLLABORATIVE MARKETING, LOS ALTOS, CALIF.

Apple to take a bite out of Mac E-mail market

BY KEVIN FOGARTY

Los Altos, Calif.

User dissatisfaction with the performance of third-party Macintosh-based electronic mail systems combined with improved messaging capabilities built into the Macintosh operating system could send droves of users to Apple Computer, Inc. for their E-mail.

That's the conclusion reached by market research firms Collaborative Marketing, based here, and Indigo Systems of San Francisco following interviews with more than 300 E-mail administrators for a study released last week.

Half of the respondents said they are unsatisfied with their Macintosh-based E-mail packages, and 35% plan to change vendors during the next two years, the report finds.

Many Macintosh E-mail administrators said they are disappointed in their third-party messaging products, but the study says vendors of those products will still grow over the next three years.

However, the growing messaging abilities of Macintosh and personal computer operating systems, as well as interoperability between these messaging technologies, will prompt many users to turn to systems vendors for E-mail.

The study predicts Apple will outpace third-party providers during the next three years, thanks to its Apple Open Collaboration Environment (AOCE). Announced last year, AOCE is a set of application programming interfaces (API) that Apple intends to use to unify third-party voice mail, fax, E-mail and remote paging into an integrated communications package using an Apple client called PowerTalk.

There is nothing inherently wrong with the Macintosh that causes E-mail systems based on it to be weaker than PC applications, said Steve York, senior manager of enterprise computing at GM Hughes Electronics in Los Angeles. But there are fewer choices in the Macintosh market, said York, who uses Microsoft Mail for AppleTalk for 7,000 Macintosh users and cc:Mail from Lotus Development Corp. for another 2,000.

Once Apple integrates E-mail into its operating system, the company is expected to underprice and outgrow more established Macintosh E-mail vendors.

Harvard teaches EDI a new lesson

University builds interactive messaging nets to link users across campuses.

BY ADAM GAFFIN

Cambridge, Mass.

When it comes to computers, nobody tells departments at Harvard University what to do.

As a result, the school's Office for Information Technology is faced with a problem that is more than just academic: How to give users of incompatible desktop platforms and networks across several campuses on-line access to personnel and financial data.

"Harvard is very decentralized, and everybody makes independent decisions with regards to computing," said Michael Laing, consultant for technology evaluation for the office.

To provide information access across its campuses, Harvard five years ago deployed an electronic messaging system called Harvard EDI that works very much like an electronic data interchange system. The system is unusual in that it is used internally, whereas EDI is typically used for interorganization or intercompany communications.

But as Harvard moves to a distributed client/server architecture, the information technology office this summer will begin turning the store-and-forward Harvard EDI system into an interactive one. A key piece of that strategy is integrating the Open Software Foundation, Inc.'s Distributed Computing Environment (DCE) into the Harvard EDI system as a framework for point-to-point communications between desktop clients and back-end servers, Laing said.

SAVING TIME

The goal is to cut end-user response time to a couple of seconds at the most, Laing said. The current system, with its reliance on gateways and store-and-forward techniques, can take as long as 30 minutes to deliver information to end users, he said.

The office built Harvard EDI on an electronic mail foundation because E-mail was the one function all of the university's disparate platforms and heterogenous networks had in common.

"In our blissful ignorance, we invented our own [EDI] protocol," despite the fact that there are numerous EDI standards in existence, Laing said.
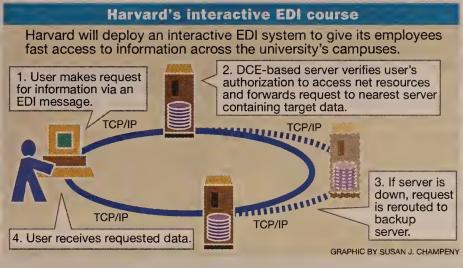
DCE was chosen to help update the messaging system because it can support a variety of network operating systems while providing standardized back-end services such as security and directories, Laing said. The DCE-based make-over will also let the university move the EDI net to Transmission Control Protocol/Internet Protocol, whereas it now uses a hodgepodge of protocols.

And because it is a distributed system, individual schools and departments within Harvard will be able to develop their own security and data viewing policies, he said. For example, the School of Medicine will be given access to all relevant records on servers across the university. But then, within the school itself, net managers can determine which workers have what sort of access to specific types of medicine-related records, as well as how they will view the information on their screens, Laing said.

Harvard's interactive EDI course

Harvard will deploy an interactive EDI system to give its employees fast access to information across the university's campuses.

1. User makes request for information via an EDI message.

2. DCE-based server verifies user's authorization to access net resources and forwards request to nearest server containing target data.

TCP/IP   TCP/IP

3. If server is down, request is rerouted to backup server.

TCP/IP   TCP/IP

4. User receives requested data.

GRAPHIC BY SUSAN J. CHAMPENY

BRIEFS

Interleaf, Inc. of Waltham, Mass., last week announced a new version of its WorldView **document distribution and retrieval** software.

WorldView 2.0 includes an application program interface for incorporating the software into other applications, support for the Standard Generalized Markup Language, a graphical browser, and enhanced displays and speed.

WorldView clients run on Windows, Windows NT, Open VMS and several Unix platforms, with pricing starting at $195 per license. The server software is available for Sun Microsystems, Inc.'s Solaris and Digital Equipment Corp.'s OSF/1, with versions planned for Windows, Windows NT and other Unix platforms this year. Server pricing starts at $9,995 for Unix versions. The Windows server will start at $4,995. Interleaf: (617) 290-0710.

LiveWorks, Inc. of San Jose, Calif., has released a new version of its LiveBoard computer **whiteboard** product that has more features and costs less. The new version includes better pen control, higher screen resolution and a software bridge that runs on a server to allow as many as 31 geographically separate locations to log on to one conference session simultaneously. LiveWorks, a subsidiary of Xerox Corp., is selling LiveBoard 3.0 for $34,900, a $15,000 reduction over the price of Version 2.0.

The product focuses around a 67-inch giant rear projection screen that acts as an interface for an Intel Corp. 486-based personal computer running LiveWorks' MeetingBoard whiteboard product. It connects via modem or any switched digital service, and across Transmission Control Protocol/Internet Protocol and Network Basic I/O System networks to users with other LiveBoard units, or on PCs. It also integrates with videoconferencing equipment from PictureTel Corp. in Danvers, Mass. LiveWorks: (408) 324-2200.

The **Object Management Group (OMG)** of Framingham, Mass., last week issued requests for proposal for a variety of back-end services for its **distributed object** architecture.

OMG is seeking proposals from companies interested in releasing public software for tracking object licensing, for keeping track of object properties or attributes and for developing an SQL-like interface for databases. All are meant to comply with the group's Common Object Request Broker Architecture. Submittals are due by July 25.

In addition, the group last week released a white paper on security within a distributed object architecture. Free copies can be obtained from the OMG. OMG: (508) 820-4300.

Desktop Data, Inc. of Waltham, Mass., recently announced a new version of its newswire utility for Lotus Development Corp.'s **Notes groupware**.

NewsEdge/Notes 2.0, will let Notes users build their own news tracking profiles by selecting from menus of more than 10,000 predefined categories.

The new version is scheduled to ship in July. Pricing depends on the number of seats and news sources selected. A typical 500-seat license for one year would cost $55,000.

Desktop Data: (617) 890-0042.

Gupta chases after Unix, 32-bit database markets

BY BARB COLE

Palm Desert, Calif.

Gupta Corp. affirmed its commitment to departmental client/server computing at its developers conference here last week by announcing new versions of its database server for 32-bit operating systems.

Database description	
Product:	SQLBase 5.2
Operating system support:	Windows NT, OS/2 2.X, NetWare 4.X and UnixWare
Pricing:	$995 for 5 users and $9,995 for unlimited users
Availability:	July
New features:	Native Open Database Connectivity support, shared cursors, enhanced integrity and more detailed error messages

Gupta will ship SQLBase 5.2 in July on Windows NT, OS/2 2.X, NetWare 4.X and UnixWare. A Solaris version will ship in early 1995. Previously, SQLBase was only available on 16-bit platforms. The firm also plans to port SQLBase and its SQLWindows application development tool to Digital Equipment Corp.'s Alpha AXP, Gupta told attendees of its International Developers Conference. Meanwhile, Gupta said it will extend its current server pricing — $995 for five users and $9,995 for unlimited users — to all database servers, according to Matt Miller, senior director of product marketing at the Menlo Park, Calif., company.

"Until now, SQLBase was a PC-centric database server. Now we're going after the low-end Unix market," Miller said.

UNIX APPEAL

The 32-bit versions of SQLBase will offer significant performance improvements over the company's 16-bit versions, Miller said. SQLBase 5.2 will appeal to the Unix market because it's a low-maintenance database server that doesn't require a full-time database administrator, yet can perform well enough to support groups of around 50 users.

Despite its aggressive stance on Unix, Gupta is not trying to compete with Unix database server market leaders Oracle Corp. and Sybase, Inc. Rather, Gupta is offering a low-cost alternative to companies that need to set up departmental Unix servers that will likely coexist with other vendors' database systems.

"SQLBase is [easier to use] and less expensive than Oracle and Sybase databases, so [Gupta has] a good chance at succeeding in the departmental server business," according to John Faig, senior research analyst at META Group, Inc. in Westport, Conn.

With SQLBase's price tag of less than $200 per user, Faig believes that Gupta will face its stiffest competition from Microsoft Corp.'s SQL Server for NT.

©Gupta: (415) 321-9500.

Harvard
Continued from page 35

Once the system is in place, an end user will be able to quickly access information stored in the data warehouse, which will be moved from the mainframe to a Hewlett-Packard Co. HP-UX server running a Sybase, Inc. System 10 database. When an end user requests information, the request will be checked by a DCE security server to see if the user has the appropriate clearance. Assuming the user does, the request will go to the nearest server housing the data — local caching of commonly used data will reduce loads on wide-area networks and the back-end master databases.

Users will also be able to quickly update databases. Laing said the data technology office is still looking at the exact method to distribute new data across a net, but Sybase's Replication Server is one possible candidate.

At the same time, departments will continue to be able to use the old EDI system to gain access to data. Laing said his office can only entice, not force, users to switch.⊡

Apple
Continued from page 35

"You can have a whiz-bang Mac mail system that works just fine or not — the underlying architecture of the Mac has nothing to do with it," York said.

But for Macintosh E-mail to become ubiquitous, it has to be part of the operating system, as it has been in Unix systems, said John Katsaros, president of Collaborative Marketing. If Apple includes PowerTalk in future versions of its System 7 operating system — including a release this summer — and Microsoft Corp. includes a Microsoft Mail client in its upcoming Windows 4.0 release, E-mail will become ubiquitous in both Windows and Macintosh environments, he said.

Add to that the Microsoft-Apple interoperability agreement to link AOCE clients and servers with Microsoft Mail clients and servers on Macintoshes and PCs via Microsoft's Messaging Application Programming Interface, and you have a combination that is hard to beat, Katsaros said.

Apple will take the lion's share of the Macintosh-based E-mail market, with current market leader CE Software, Inc. falling to second place, according to the study (see graphic, page 35). The freeware version of Eudora available on the Internet and the for-fee version from Qualcomm, Inc. of San Diego will also draw adherents.

York has no plans involving AOCE. He does plan to phase out Microsoft Mail as the firm discontinues development of its AppleTalk mail server. Microsoft's plan to convert its Macintosh users to Touchdown, which will run only on a Windows NT server, "leaves us kind of cold," York said.

©Collaborative Marketing: (415) 949-3256.

D&B makes smart move, budgets new development and mgmt. tool

BY KEVIN FOGARTY

Dun & Bradstreet Software has released a new budget development and management module as the latest addition to its Smart-Stream series of client/server applications.

SmartStream Budget is built on top of Microsoft Corp.'s Excel 5.0 spreadsheet, which it uses for a front end, and Sybase, Inc.'s SQL Server back end. It works with Smart-Stream decision support, work flow and reporting tools from D&B Software.

The application lets users put together their budgets in a popular spreadsheet, rather than forcing them to use an unfamiliar and hard-to-use proprietary interface, said Doug Gosling, manager of financial systems at Toronto-based Ontario Hydro.

"The beauty of it is that it integrates with something you already have — your spreadsheet platform," Gosling said. "And the budget process is basically spreadsheet-based, so this makes budgeting just another option on your menu. And the infrastructure [D&B] built on top of Excel lets you import real data and control administration and security."

In its first iteration, released last week, SmartStream Budget runs on Hewlett-Packard Co.'s HP-UX and Data General Corp.'s DG/UX servers. Next month D&B Software promises to release versions for Sun Microsystems, Inc.'s Solaris and IBM's AIX.

INTERNAL INTEGRATION

The product is designed to integrate with other SmartStream products such as Financial Stream, which could be used to automatically pull out data on the previous year's actual spending and boost every category by 10%, said Dan Drechsel, director of program management for D&B Software. SmartStream Budget will also integrate with SmartStream Decision Support and HR Stream to let users do what-if projections on head count and salary levels.

Integration with other SmartStream products and using Excel as a front end are advantages, but the whole SmartStream product line would be more attractive if it ran on more than just Sybase's database, said Bob Marthinsen, program manager for information technology planning and strategy for Alcatel Network Services, Inc., a telecommunications switch manufacturer in Richardson, Texas.

The SmartStream line is built using SQL Server as a data repository and to handle stored procedures, Drechsel said. SQL Server is packaged with SmartStream applications specifically to be used with those applications, not as a general-purpose database.

"We've been bugging [D&B Software] about that because we've got a lot of Oracle in-house," Ontario Hydro's Gosling said. Although Sybase gateways as well as D&B Software application development tools and InterQ middleware let him move data between SmartStream and other databases, it would be easier if he could choose the database, said Gosling, who has already bought the budget application.

> "The beauty of it is that [SmartStream] integrates with something you already have — your spreadsheet platform."

SmartStream Budget snapshot

Operating system:	HP-UX and DG/UX.
Database:	Sybase SQL Server built into application.
Cost:	Starts at $60,000.
Integration:	Works with D&B Software's SmartStream Decision Support for reports, queries, analysis and report distribution.
Still to come:	Versions for Solaris and AIX. Both will be available later this month.

Although D&B Software plans to port SmartStream to other databases, the company is concentrating more fully on extending the product line. Drechsel had no projection for when new database ports will be available.

Overall, the product looks easy to use and is powerful enough to meet the budget needs of most users, said Jennifer Scholze, an analyst at International Data Corp. in Framingham, Mass.

Gosling agreed, saying SmartStream Budget's GUI and spreadsheet interfaces make it easier for his users to work on the system. With his previous character-based budget application, "lots of people were doing budgets in Lotus [1-2-3] or Excel and importing them into the system," he said. ⊠

Client/server firms release application development tools

BY ADAM GAFFIN

Several client/server application development tool vendors, including Intersolv, Inc., Bachman Information Systems, Inc. and Uniface Corp., last week announced new versions or enhancements to their software.

A recent spate of tools announcements could presage a shakeout in the market over the next year as users learn to separate out vendor claims, one analyst said.

"Each of the vendors now is trying to pretend they do everything," said Ed Acly, director of software research at International Data Corp., a market research firm in Framingham, Mass. "The next 12 months are going to be 12 months of revelation."

Intersolv, Inc. of Rockville, Md., released a new version of its Excelerator II 2.0 tool set that adds support for developers using Windows and Windows NT platforms. The announcement was made at the company's annual user conference in Washington, D.C.

Excelerator II is an object-oriented modeling tool that lets developers use a graphical interface to design and then compile applications. The new version adds better tools for documenting and mapping networked and database applications and adds support for Oracle Corp.'s Oracle6 and Oracle7 databases.

The software's LAN Repository, used by developers to share code, can already be built atop databases from IBM, Microsoft Corp. and Sybase, Inc. It supports several communications protocols, including Advanced Program-to-Program Communications and Transmission Control Protocol/Internet Protocol.

Acly said Excelerator II represents a good effort toward extending the power of traditional computer-assisted software engineering, particularly data modeling, toward the development of distributed applications.

Like many other object-oriented tool vendors, Intersolv said last week that it does not offer support for the Object Management Group's Common Object Request Broker Architecture (CORBA) for distributed object computing.

Retooling		
Product	**Price**	**Availability**
Intersolv's Excelerator II 2.0	$4,000 per license	Now
CSA SilverRun/ Uniface bridge	$2,500	Now
Bachman's Ellipse	$20,000 for Unix repository, $1,999 per user	Third quarter

Mike Merriman, a product manager at Intersolv, said it makes little sense to do so because the current version of CORBA does not include specifications for linking object request brokers from various vendors.

CORBA 2.0, which is expected later this year, will provide such specifications. Merriman said that adding support for CORBA 2.0 would be relatively painless for Intersolv, given its existing rules-based, object-oriented architecture.

Bachman, meanwhile, announced a new tool called Ellipse for building and deploying client/server applications across an enterprise. It features automatic reestablishment of communications for downed servers, automatic application partitioning after the application is installed, and a variety of version-control, load balancing, management and interface development tools.

Set to ship in the third quarter, Ellipse will support Microsoft Windows and OS/2 for Windows.

It comes with a runtime version of Sybase Inc.'s SQL Server, and can talk to IBM DB2 and other databases through Sybase's OmniSQL Server middleware. It will also support TCP/IP and Novell, Inc. NetWare networks.

Bachman plans to add support for other vendors' graphical development tools, such as those sold by Powersoft Corp. and Microsoft.

Also last week, Uniface Corp. of Alameda, Calif., and Computer Systems Advisors, Inc. of Woodcliff Lake, N.J., announced a bridge between their Uniface Six and SilverRun tools that lets developers working with either product export changes made in one to the other.

This will give developers a way to import SilverRun data models into Uniface Six's object repository and to automatically update SilverRun models when making changes in Uniface Six applications.

SilverRun is a development workbench that supports Windows, OS/2 and Macintosh applications. Uniface Six, intended for client/sever application development, supports a variety of platforms and back-end databases.

©Bachman: (617) 273-9003; CSA: (800) 537-4262; Intersolv: (301) 230-3200.

WE UNDERSTAND THE VALUE OF TIME.

We Swiss are preoccupied with time. We measure it. We monitor it. And we make every minute count.

No wonder that at Swiss Telecom, we are committed to make your European network run like clockwork. Making us your primary or secondary hub is the best conceivable way to assure that your calls, data and video are where they need to be, when they need to be there—without unnecessary detours, without delays and without excuses. What else would you expect from the world's timekeeper?

But there are other ways our penchant for punctuality becomes your gain. At the first sign of a problem, for example. That's when you'll learn the value of a Leaseline Control Center that monitors your circuits continually and is staffed with English-speaking specialists 24 hours a day, working closely with you to take care of little problems before they become big ones. Or when you have a question about a PTT's regulations, an overseas tariff, or a possible change in your network design. That's when our Washington, D.C.–based staff is ready to respond to your requests on a moment's notice.

What's punctuality without precision? We provide full connectivity through high-quality, high-speed digital gateways—fiber, satellite, and radio links offering all the capacity you'll ever need. And made easier by our partnership in Unisource, a pan-European carrier with global reach.

Call Swiss Telecom North America at 1-800-966-1145 to find out more. We urge you to do so today. For if your business depends on European or worldwide telecommunications, you haven't a minute to lose.

SWISS TELECOM ᴾᵀᵀ⊹

Get smart on sales

"IBM...must realize the need to simplify their approach — minimize the complexity and work on educating the sales staff."
Keith Weiskittle, Miami Valley Hospital.

As I reviewed this week's Special Issue on IBM networking, Weiskittle's comment really stuck in my mind because it voiced a concern I have heard from a number of customers.

Many users feel that vendors aren't training their sales forces to understand the complex technologies they are pushing and, more importantly, how those technologies can be applied to help individual companies. They complain that sales personnel don't even understand what their companies sell.

This issue looms larger as the industry evolves. It's not a problem with companies that sell a single product. But as companies merge, expand their lines and branch into new technologies, things get sticky.

Consider the LAN internetworking arena. You have hub and router companies weaving in switching and routing while rolling out ATM support. To make the transition successfully, sales forces not only have to understand how these new technologies work, but how they interoperate and how customers will migrate to them.

Or take Novell, whose resellers have to deal with a product line — NetWare 3.X and 4.X, AppWare, UnixWare and, now, the WordPerfect and Borland products — that has gotten far more diverse in just a couple of years.

This issue is particularly important for broad-based suppliers such as IBM, which sells virtually everything. Sales reps have to understand the LAN strategy, the network management plan, the ATM migration scheme, the new Any-Net technology and more — and that's just the physical network side.

IBM has a lot of great technology on the table or under development, but it's not clear that the sales force can communicate that. (Consider Weiskittle's struggle to learn more about IBM's configuration, installation and distribution products, described in the Q&A on page 49.)

And it isn't just the sales reps who struggle. It is not uncommon for top IBM officials to be asked during strategy briefings: "Why doesn't anybody know about this product?"

IBM had considered breaking up into smaller companies — Baby Blues — partly to improve sales efficiency. But it has wisely avoided that course (customers don't want more sales reps) and is aligning its sales forces with key vertical industries. That should improve understanding of customer needs.

But more important for IBM and other companies is a commitment to improving sales training. Customers want to build "partnerships" with suppliers. But in this time of tight budgets and constrained resources, buyers don't have the luxury of hand-holding the sales force.

JOHN GALLANT   jgallant@world.std.com

TELETOONS FRANK AND TROISE

Look at stratum 17, professor! The latter half of the 20th century... a sudden disappearance of Studebakers, wide lapels, mood rings and ISDN...

Industry shakeout won't leave net users in the rubble

In many ways, our industry is similar to seismology, the study of the earth's tremors. Lately, we've seen a lot of turmoil among the hardware, software and networking companies upon which we build our information systems (IS) foundations.

Like two of the earth's great plates coming together, companies like Novell, Inc. and WordPerfect Corp. unite with a bang, sending customers and competitors reeling for awhile.

But unlike earthquakes and other natural disasters, this recent industry shakeout actually has some beneficial effects. I, for one, am pleased about the Novell/WordPerfect merger and some of the other recent unions, such as Symantec Corp. and Central Point Software, Inc.; Banyan Systems, Inc. and Beyond, Inc.; and, Artisoft, Inc. and Eagle Technology, Inc.

I used to be negative about computer company mergers, thinking that they weren't necessary in order for two companies to produce products that interact well together. But I've changed my tune now that I've done some research on the cost of support and integration of these supposedly compatible products we've all been buying.

In a rather well-known 1992 Nolan, Norton & Co. study on managing end-user computing, the research firm concluded that the cost of support of these end-user computing systems (mostly PCs and LANs) is much higher than anyone suspects. Nolan, Norton contends that support costs might be as high as $8,000 to $22,000 annually per PC or workstation.

It's my theory that a significant portion of these support costs comes from just trying to get two (or more) products to work together. You know how much time you spend trying to optimize Windows under NetWare or to get the gateway working between two different electronic mail products. While I have no empirical evidence, I suspect that better product integration can help reduce these outrageous support costs.

So what does this have to do with the benefits of computer company mergers? Plenty. I see this whole trend in industry consolidation as one big push toward complete, integrated solutions. We're moving toward an industry that will be just like the old days, when two or three dominating vendors offered total integrated computing solutions.

Let's look at Novell and WordPerfect as our example. Before the announced merger, Novell was dominant in one area — desktop networking — yet weak in offering the integrated workgroup applications that people need. (Even the Nolan, Norton study "verified the importance of new integrated capabilities that go beyond individual work, individual packages and individual locations.")

As for WordPerfect, the firm was largely known for its top-rated word processing products, but I believe the real value of the company was in its workgroup application WordPerfect Office. This one, integrated product offers more than messaging; it provides work flow, calendaring and group scheduling, as well. It even has a telephone interface for your E-mail and calendar. Add WordPerfect InForms to the picture and you have a powerful combination for managing the flow of data. But still, something was missing.

That's why WordPerfect acquired SoftSolutions Technology Corp. earlier this year. SoftSolutions has a nifty product for complete document management, including full text searches, and document retrieval and archiving.

Now, with WordPerfect and SoftSolutions becoming part of Novell, we're starting to make some sense of this workgroup computing thing. Novell brings the underlying infrastructure — the means to tie everything together — to the table. WordPerfect brings the way to create and disseminate documents, and SoftSolutions brings the means to manage them. That's starting to look like a tasty offering to many IS managers I know.

Of course, it will take a year or more to integrate all of these disparate parts. Even new Novell CEO Robert Frankenberg says his first order of business is to transform Novell from a network operating system company to a "total package supplier" company. When that happens, the new and improved Novell will have a mighty powerful story to tell. Why, it might even be able to stand up to its arch-nemesis, Microsoft Corp.

Some people might not like the idea of just two or three major solution providers dictating what applications run on what net platform. I'd argue that point. The more I work with clients who are trying to make sense of this chaotic microcomputer industry, the more I believe that IS people don't want to have to make decisions. I believe they just want to be told what to buy. Evaluating products takes time, and there's too much opportunity for picking the wrong products. That could cost someone his or her job. Remember the old adage, "No one ever got fired for buying IBM?" I think soon we'll substitute Microsoft or Novell or Lotus in place of IBM in that statement.

Oh, yes, I can see a downside to all of this consolidation. The little guys — the cottage industries that lack the marketing might of Microsoft — will lose out. Even once big companies like Borland International, Inc. and Software Publishing Corp. are feeling the pinch. We'll see a loss of "best of breed" software as buyers have to settle for what's available from the industry powerhouses. We'll lose innovation and our right to choose. (A survey published by International Data Corp. in February shows that IS managers are beginning to sour on suites — discounted bundles of software applications — with 55% of those surveyed expressing distaste.)

But we'll get over it, because in the long run, the tighter product integration will lessen our burden of support, saving us time and money. The truly successful mergers will prove to be those that take disparate product lines and make them work better together to meet real user needs. I think that's where Borland failed in its merger with Ashton-Tate Corp. a few years ago. The product lines were never successfully brought together, which divided the market and drove customers into Microsoft's waiting arms.

Before all is said and done, I think we'll have a lot more tremors in the network and computing industry. And who knows, we just might see The Big One soon.

Musthaler is vice president of research at Currid & Co., a Houston-based technology consulting firm. She can be reached at (713) 789-5995.

Opinions

SECURITY PERSPECTIVES
by Michel Kabay

Virus contest isn't just harmless fun

An organization that thinks viruses are an expression of free speech is currently sponsoring a virus-writing contest. Its example of an amusing new virus is one that interferes with free speech. Crazy? Probably. Evil? Definitely.

It's time to turn our wrath on the sociopathic scum who deliberately write viruses and encourage others, mostly juveniles with inadequate personalities, to write malicious software. These people are costing user businesses untold amounts of time and money to repair the damage viruses wreak on networks.

Viruses are free-ranging, self-reproducing code. They attach themselves to your software and then use your computer resources to copy themselves into other executables, possibly damaging data and even sometimes causing your network to become overloaded and crash. It's hard to count viruses because experts disagree on which ones are merely variations of one another and which ones are distinct. Nonetheless, the antivirus industry estimates the number of PC viruses to be in the thousands.

The organization sponsoring the current virus-writing contest, slated to culminate at the end of this month, is a U.S. publisher of virus source code. The purpose of the contest is to reward the person who can come up with the most innovative virus.

One example they suggest is a virus that generates junk faxes. Another suggestion is an automatic censor program. This hypothetical virus would stop your word processor from using words on the virus' list of forbidden terms. To start with, it might insist on politically correct, gender-neutral writing. However,

unlike utilities that check grammar and style, the anti-free speech virus would block all attempts to use forbidden terms, regardless of context. The virus author could make the virus flip to a different list after a month or two; the anti-free speech virus would then force people to use politically incorrect terms such offensive language.

Suppose amateur, genetic engineers were invited to submit new, biological viruses. What if someone offered a prize for a new drug with which to contaminate public water supplies? There would be an outcry of rage against the organizers of such a contest. They would be charged with reckless endangerment — acting deliberately in a way that any reasonable person would agree would likely cause harm to others.

This virus-writing contest is reckless endangerment.

The mere presence of unauthorized software threatens the confidence with which we use our nets. Even a virus intended to be harmless causes careful users to spend time verifying data integrity on infected systems. Anyway, virus authors do not apply proper quality assurance to their creations; they lack the variety of hardware and software — let alone network configurations — to test their malicious software. Virus authors cannot ensure that their products will not cause data corruption on your servers or halt your network operating system.

How can we distinguish between legitimate antivirus researchers who need to tinker with viruses as they tune their products and irresponsible lowlifes who play with these things for fun? Let's turn to biological virology for ideas.

Keeping disease-causing organisms (pathogens) in a laboratory is strictly regulated. Playing

with biological viruses is not a question of First Amendment rights. Standards of containment limit distribution of virulent strains; violation of the standards leads to severe criminal penalties.

Pathogens are attenuated to limit their effects; similarly, experimental computer viruses could be written with controlled payloads that would identify the viruses but prevent them from doing any damage.

Micro organisms and biological viruses can be tagged with unique genetic markers; similarly, computer viruses created by honest antivirus laboratories could be marked with cryptographic authentication codes.

Experimental pathogens are often incapable of reproducing outside a controlled environment; such enzyme-deficient organisms die without rare nutrients. Similarly, computer viruses written by honest people could be made dependent on special codes placed on disk or in memory by their authors but never found in the wild (that is, on ordinary systems and networks).

If a laboratory computer virus was to get out and cause harm to anyone, victims should sue the author of the virus in civil court. Someone publishing viral source code should be sued if any of his published viruses ever appeared on anybody's computer or network.

Don't be fooled: Writing a virus is not an exercise of free speech. It's mischief.

Writing uncontrolled, untraceable viruses is evil and irresponsible. I urge network and system managers to hold virus publishers responsible for the consequences of their actions and push for injunctions to stop their dangerous contests.

➤ Kabay is director of education with the National Computer Security Association in Carlisle, Pa. He can be reached on the Internet at 75300.3232@compuserve.com or by phone at (514) 931-6187.

Letters

Virtual LAN fan

I found your special feature on virtual LANs (May 2, page 1) to be very interesting. I think you have investigated several major players in the switched hub arena, but I would like to point out one you made no mention of.

Hughes LAN Systems, Inc. has a hub selling now that is easily as impressive as Cabletron Systems, Inc.'s or SynOptics Communications, Inc.'s announced hubs. I have been researching hubs for more than a year, and I am convinced that Hughes has the best hub on the market today.

Also, I believe a big problem with switched nets, in general, is in their handling of multicast. You alluded to this in your article's list of specifics to seek from vendors. I

would have liked to have seen you address this in greater detail.

There is a protocol that many router vendors are beginning to offer for Internet Protocol multicast propagation, known as the Internet Group Management Protocol.

This basically allows a node to join a multicast group across router boundaries. This same type of protocol could be used with virtual LANs given routing capability (or just multicast gateway capability) in the hub.

However, it seems to me that in a switched environment, a better mechanism would be to allow nodes to dynamically join multiple virtual LANs.

This way, the entire virtual LAN would not have to receive data that is uninteresting to the majority of nodes. In other words, there could be a set of virtual LANs used for

multicast-like functionality that would be dynamically allocated.

This would also let you join a group that is, for example, participating in a particular videoconference while still retaining connectivity to other multicast (virtual LAN) groups.

All of the hub vendors I have examined only allow static assignment of virtual LANs and single membership. Hughes' hub uses Hewlett-Packard Co.'s OpenView to assign virtual LANs. This requires intervention from the systems administrator.

It would be useful if there was a protocol or extension to an existing protocol that would allow the dynamic joining of any particular node to a set of virtual LANs delineated for such multicast-like use and the hubs were intelligent enough to facilitate it. Are you aware of any work being done on such a beast?

Matt MacPherson
Network administrator,
Accelerator Division
Fermi National Accelerator
Laboratory
Batavia, Ill.

John Morency, co-writer of the virtual LANs feature, replies: You are correct in stating that most, if not all, virtual LAN vendors use static means to assign end stations to virtual LANs. This is due to the early stage of both virtual LAN technology and its associated management services.

The protocol you propose would require a proper set of both configuration and security management services in order to obtain mainstream industry acceptance.

The definition and construction of such services is a bit beyond where most vendors are today. As such, we are not currently aware of any mainstream hub or router vendor who is actively developing either the protocol you cite or any related management services.
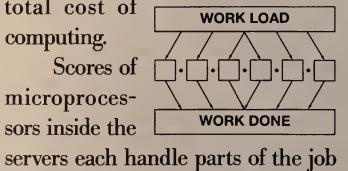
Parallel processing is almost as amazing as Nature itself. And IBM is leading the way with parallel servers for both the System/390® and UNIX® environments.

Like Nature's most efficient creatures, IBM's parallel servers divide up tasks to accomplish bigger jobs quickly. And this greater efficiency significantly reduces the total cost of computing.
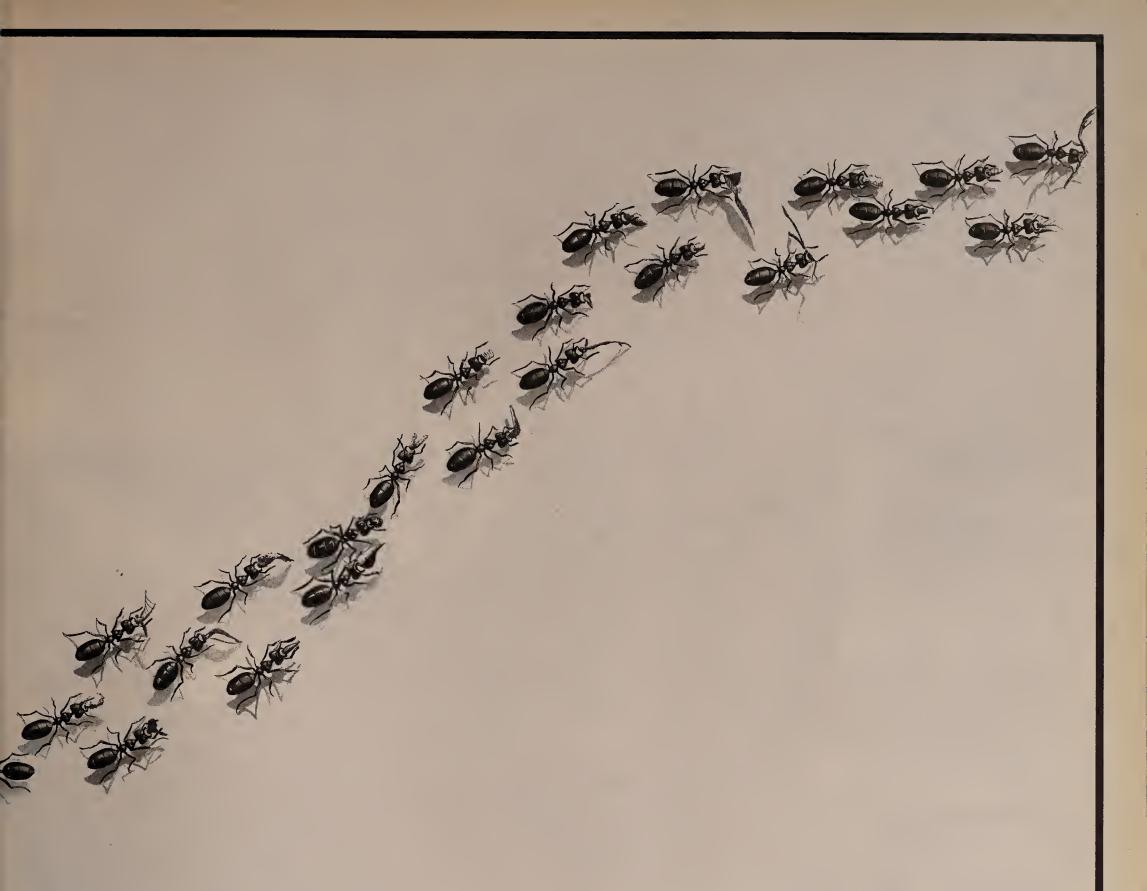
Scores of microprocessors inside the servers each handle parts of the job

WORK LOAD

WORK DONE

(just like these little guys), while the system as a whole coordinates and seamlessly completes the project. Everything from existing applications to emerging ones like data mining and video-on-demand.

And recent, dramatic IBM breakthroughs are what have put

Announcing IBM's new parallel processing.
Its division of labor conquers bigger jobs for smaller costs.

it all within your reach. In the System/390 environment, IBM's exclusive new S/390 Parallel Sysplex technology is the key—adding powerful data-sharing capability to your existing system. And in the UNIX environment, the new IBM POWERparallel™ SP2™ now makes both

numeric-intensive *and* data-analysis processing practical for your system.

Whichever environment you use, your current investment is protected—and, indeed, greatly enhanced.

Only IBM has the innovation, the resources and the commitment

to customer service to offer you this affordable, unparalleled choice in parallel technology.

So join the ranks of IBM parallel processing today. By dividing the labor, you'll conquer the costs.

IBM®

End the long wait
for reconnecting sessions

VTAM

Bringing back sessions following an application or network failure takes too long. In fact, it can take longer than the actual recovery. And the more terminals on your network, the longer it takes.

But not any more. BMC Software, Inc.'s new SESSION PERSIST re-establishes sessions fast.

When the session reconnect is fast, your entire recovery is faster. So, you can end the problem of help desks swamped with end-user calls, get critical tasks moving and minimize money lost due to application down time.

SESSION PERSIST also dramatically cuts related resource problems: storage spikes, CPU usage and network traffic.

Full availability in record time is here now. To learn more about SESSION PERSIST, call BMC Software at 713 918-8800 or **800 295-4BMC**.

IBM Business Partner

BMC® SOFTWARE

The Experience. The Technology. The Future.

Users give IBM an earful

DAVE CUTLER

Everywhere you turn today, IBM execs will tell you they're taking their cue from customers. It sounds great, makes good press. In the end, IBM's success will hinge not so much on listening to customers as it will on the vendor's ability to address user needs.

Consider this insight about IBM's net management tools from Stephen Hall, senior vice president of information systems at MedAmerica Health Systems Corp. and a member of IBM's Networking Systems Customer Council.

"Our advice is to offer a smaller suite of products rich in function — even if they're optional functions. Reducing that confusion factor will allow customers to take steps. [IBM] has confused customers so much, many aren't doing much in this area."

Sanjiv Ahuja, IBM's director of enterprise management platforms, says IBM always weeds out unsuccessful products from its line. Maybe so, but the issue isn't about too many products; it's really about the need to address confusion in the customer ranks. IBM needs to simplify its many messages to users.

The extent to which customers pledge allegiance to IBM in future years will hang largely on how well users understand IBM's vision, and equally on how well IBM listens to its customer base.

We're eager to listen, too. We want to hear what you need from IBM and other SNA providers. What are you asking for that you're not getting? How has IBM and other vendors responded to your SNA needs? Drop us an E-mail; let us know what you think.

➥ **CHARLES BRUNO** cbruno@world.std.com

JEFF SCHEID

Trio of users open up

and discuss the evolution of their enterprise nets, identifying short- and long-term needs for IBM to address. *Page 49.*

Ready to route SNA?

Not so fast. Router vendors have formed a myriad of techniques for routing SNA traffic on LAN internets. There are several forms of IP encapsulation, Cross-Comm offers PIR, and IBM served up DLSw. Which is right for your net? Maybe you should consider SNA routing. *Page 53.*

APPN grows up

To technical zealots, IBM's Advanced Peer-to-Peer Networking and the Transmission Control Protocol/Internet Protocol have about as much in common as OS/2 and Unix. Yet APPN is evolving to offer many of the same services as TCP/IP, and then some. *Page 65.*

Staying in control

IBM net exec talks about managing LANs on an enterprise scope, mapping out the role of NetView for centralized and decentralized control. *Page 69.*

JEFF SCHEID

Users pipe up

> **Trio of net execs say IBM is taking its cue from users, but it still needs to simplify product line to stem confusion.**

In keeping with its theme of listening to customers, IBM executives in Las Vegas gathered at the recent NetWorld + Interop conference with eight customers who form the IBM Networking Systems Customer Council.

The council meets with IBM periodically to review tactical plans and strategic directions for the vendor's products and services. The council also gives IBM the opportunity to gather customer requirements, understand buyers' investment strategies and strengthen product offerings based on customer feedback.

Three members of the Networking Systems Customer Council met at Interop with Network World editors to discuss their current and future product needs, explore the issues they are grappling with and assess how IBM is changing to address their concerns.

Daniel Grange, director of information technologies at PSA Peugeot Citroen, the Paris-based automaker; Stephen Hall, senior vice president of information systems at MedAmerica Health Systems Corp.; and Keith Weiskittle, director of computer operations and technology at Miami Valley Hospital in Dayton, Ohio, shared their thoughts with NW Editors Charles Bruno, Michael Cooney and Paul Desmond.

IBM's Systems Network Architecture turns 20 this year. How is your company using SNA, and do you see it playing a pivotal role in your network infrastructure five or 10 years down the road?

Hall: We're a little different from some IBM customers in that our IBM processing only plays one part of our overall computing strategy. Our focus has always been around the network as the computing facility. SNA just happens to be one protocol that runs across our network that serves one subset of our customers in terms of gaining access to computing capabilities. To answer the question, SNA will continue to be a mainstay; it will probably continue to provide the main backbone links into our IBM processors. We do that through token-ring networks, predominantly. And we don't see that going away in the short run.

Weiskittle: I think the key to what Steve said is that a lot of people, a lot of big glass houses, think of the mainframe as the center with LANs branching out. We've always thought of the network as the center of our world, and it just so happens that the 3090, — and the Tandem and AS/400s — sit on the net. So when we get together with the rest of the companies in the Networking Council, we bring a different view. Network management for us has to be network-based, not host-based.

Grange: At Peugeot, we maintain an SNA network of about 20,000 terminals and all 3270-based applications are always growing. We think in the next five or 10 years, we will continue with SNA in the way we have today. But the other main activity of Peugeot is with CAD/CAM, and we have big growth in the number of workstations tied onto this type of TCP/IP network. What we want to do today is to place some gateways between the SNA and the TCP/IP networks. So the problem now is to provide the ability for the user to access Unix applications or 3270 applications from the same point. We think the two networks will merge in the future.

What is your impression of IBM's router strategy? Are routers even important to you?

Grange: The router strategy is very important for us, and we think that IBM has to be in that market but they are a bit late. And it's why we have not been able to wait for them.

Weiskittle: We've implemented IBM 6611 routers, and they're important to us as we start to reach out and connect to a lot of different health care organizations, which have a need to share data with both sides. For example, we have a connection between us and an insurance company where the intent is to allow us to check the insurance status of various customers as they come to the door. The potential there is great as we start to work with EDI and electronic medical records and that whole process in health care. I agree with Daniel [Grange]: It appears that IBM has been late into the game. Our experience in installing the 6611 was a rocky one, just because I think we felt like we were implementing at the same time that IBM was getting experience. We learned alongside them in some cases.
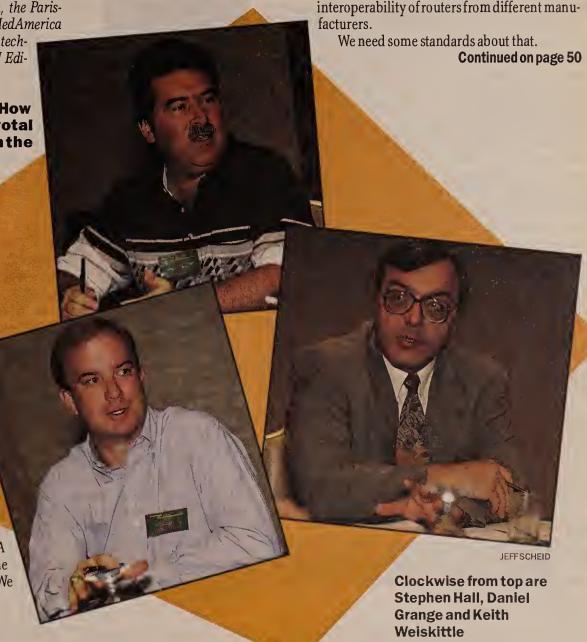
What improvements could IBM make to the 6611, or to its router strategy overall?

Hall: Some of our problems early on focused on interoperability with other vendors' routers. In our case, it was Proteon. Some of those were basic start-up issues of interoperability, but we worked out most of those.

Weiskittle: It was such a long time ago, but concerns with being able to route SNA is, of course, a big problem for us. Supposedly, Proteon and IBM had some solutions worked out that didn't ever really pan out.

It seems to be that, not only IBM, but all the router companies have to get together and try to develop some standards, something that hasn't happened so far. That seems pretty apparent to us from what we've seen.

Grange: I think IBM can focus its 6611 efforts in terms of performance improvements. We need instantaneous power. And I think the best issue is to understand the problem is interoperability of routers from different manufacturers.

We need some standards about that.

JEFF SCHEID

Clockwise from top are Stephen Hall, Daniel Grange and Keith Weiskittle

Do you have any plans to implement Advanced Peer-to-Peer Networking technology. If so, what is the added benefit you see coming from that?

Grange: Today, we see no benefits of going to APPN. The classic SNA is still very robust and very reliable. We think in the future we will go to APPN, but the technology and the routers to support it are just not yet here today.

Let's talk about network management. Are you satisfied the current crop of tools are mature enough for you to effectively manage your local-area networks?

Weiskittle: I think it's a very immature market at this point. I think IBM has a stable enough product and we're willing to go forward with what they have. But completely and totally comfortable? Probably not yet.

What do you need to manage at the LAN level?

Weiskittle: What we want to be able to manage is to have different types of configurations of workstations — be it at the department level or functional level — and to centrally distribute the operating system, standard desktop configurations and standard suite of packages to each of those different categories of network nodes.

To be able to do that from a central location is going to save us a tremendous amount of time and let us keep our staff low.

> **66** Our advice to [IBM] is to offer a smaller suite of products rich in function — even if they're optional functions. **99**
>
> **Stephen Hall**

How do you plan to manage your enterprise networks in the coming years?

Hall: Systems management, configuration management, and net management are subjects the Network Council has paid a lot of attention to over the last several meetings. I don't think there's been a well-defined strategy within IBM or many of the other vendors in terms of network or systems management. These products have grown up based on individual platforms; there has not been a well-unified strategy. It's only recently that IBM has — after hearing a lot of input from us over two years — confirmed what we have been saying. That is, there's just too many products that aren't well coordinated. They understand that now. They're identifying all their products, finding where there's commonality. I believe we'll start to see some weeding out of duplication of products and moving toward a unification of systems and net management products.

From our company perspective, we're looking for system or network management products that are looking from the network out. So we're looking for products from that focus on network functions, as opposed to being geared to one platform such as a 3090 or an AS/400.

Do you expect to use a product like a Net-View/6000 or allow the mainframe to retain more of the management role?

Hall: I would suspect our focus will not be from the mainframe. That is really not consistent with our network focus. Net-View/6000 in all likelihood will be the focus. There is a lot of IBM attention going to that product. While we are not a RS/6000- or a Unix-focused company, we will probably move in that direction because that's where a lot of the products are being developed by IBM.

What are you using now?

Weiskittle: We have various tools, like ProTools and LAN Network Manager. A lot of separate individual databases that we've maintained for managing various nodes, but nothing that brings it all together.

One of the things we've brought to the Network Council is, as a small shop, we've let them know that not everybody might be willing to jump off the boat and go for an RS/6000. They need

to maybe focus on some OS/2 management solutions. Although we won't hesitate to go to the RS/6000 if the benefits are there, probably today the most natural environment for us to implement would be an OS/2 platform. I think IBM has heard that, although their stated direction is they are going to start with products in the RS/6000 area and develop the OS/2 platform.

What steps has IBM taken to encourage you about its management direction?

Hall: There are now projects under way aimed at identifying all the network and systems management products and their functionality, so IBM can start to understand how these products overlap. IBM has stated a direction of moving from these diverse products toward fewer products with additional features. They understand they need to reduce and remove a lot of confusion associated with all their diverse groups of products.

Our advice to them is to offer a smaller suite of products rich in function — even if they're optional functions. Reducing that confusion factor will allow customers then to take steps. They have confused customers so much that a lot of customers aren't doing much in this area.

We think they'll sell more product and we'll be able to deliver more management of our networks and systems when they reduce the confusion.

How does Peugeot manage its vast net?

Grange: The route we have taken is to centralize the management. We use NetView, of course. Our problem is that we do not have enough knowledge on each side of the backbone network to manage the entire network.

So you need to bring that management into a central site?

Grange: That's right. We do that today with a central team that is managing the network 24 hours a day, seven days a week. We have two separate management programs: the SNA network is always based on NetView and NetView hosts; and we chose to use NetView/6000 to manage the TCP/IP side. What we wanted was a central station to manage all the network. In fact, it's not possible.

Today, we are on the way to implementing some remote and decentralized stations communicating with the central station based on NetView/6000.

What we plan to do is try to have a unique view — a single console — to manage all the network, but we can't avoid using network management on the host system, too. We will look to add a complementary capability between the host system and the NetView/6000 stations.

> **66** We need a unique way for users to go from a workstation to all the applications on the net. **99**
>
> **Daniel Grange**

What are you telling IBM you need for managing your networks — from the host perspective, as well as from the distributed platform?

Grange: I think we must choose the power of the host for a large amount of work. The station has to be a concentrator of problems and alerts. Our problem management, for example, is completely done on the host. I think the main problem is to manage the alerts on the net. IBM has to help us do that.

We spoke earlier about the need for an OS/2 version of NetView. But if you wind up using Net-View/6000, what would you like to see from IBM in that product that it doesn't yet offer?

Hall: Our largest problem comes from our workstation growth, and it is the ability to manage software, to distribute software, workstation configuration, asset management — those types of things that have become a larger management problem.

And since we do manage those centrally, as the network grows, it becomes a problem. We're looking for the ability, centrally, to manage all of those things.

Didn't IBM announce its Configuration, Installation and Distribution (CID) specification to address the issue of software distribution?

Weiskittle: CID is not really a product but a strategy and a group of products. We will be adding CID functionality this year. Initially, we will be using CID concepts and technology for distributing OS/2 operating systems, applications and standard desktop wares across our LANs to all the workstations. That is a product or concept that has been developed by IBM that I think works. We're going to implement it here very quickly this year.

> **66** One of the things we've brought to the Network Council is, as a small shop, we've let them know that not everybody might be willing to jump off the boat and go for an RS/6000. **99**
>
> **Keith Weiskittle**

Have you been happy with the rollout of the CID-compliant products?

Weiskittle: Well, it goes back to a bit of the confusion and the number of products and the different platforms. We can't tell you the frustration that we've had in maybe three or four meetings with IBM. We were an anxious customer coming to them saying, 'We want to do CID. Tell us what we need to do.' Three or four meetings later — after you hear about Net-View/6000, NetView host, NetView DM and a host of other products — you get blown away. Finally, after about the fourth meeting, we said forget the host, forget RS/6000s for now. We want to do CID, what do we need? In the process of that time frame, probably a year and a half, IBM refined the approach and started to weed out some things and focus on real products and the proper way to sell this. They were finally able to give us a service offering.

What are the specific types of products you would like to see from IBM that would help you better manage Peugeot's LANs?

Grange: We need a lot of automation because on a large network, it's impossible for a small team to manage everything. Some problems need to be resolved automatically. That's a key to keeping down our staff size and letting us do more with less people.

What is the most important advice you can pass on to IBM?

Grange: I think IBM really needs to address the gateway link between SNA and TCP/IP. Also, we need a unique way for users to go from a workstation to all the applications on the network.

Hall: Our advice is that we're looking for strategy from IBM; IBM needs to clarify its stance in terms of next generations of networks. I'm referring here to ATM strategy and some of the other higher speed net options. As part of the council, we've had an opportunity to hear part of IBM's strategy. I think the real issue for IBM will be to clarify its strategy, gain support from customers and for IBM to stay the course of that strategy. We want to be investing our dollars for the future. We're anxiously awaiting ATM deliverables. Others may question the need for ATM; we're ready to use it. We have a need for that type of bandwidth. We're waiting for delivery of product.

Weiskittle: I am convinced that one of IBM's greatest strengths with regards to ATM is that it can provide a solution across the entire network spectrum. A lot of vendors pitch their ATM devices and routers, but to be able to go to one company and get a solution from workstation to high end — the entire spectrum — I think is something IBM needs to focus on and sell as a strength.

The other thing IBM needs to do, in my opinion, is they must realize they need to simplify their approach — minimize the complexity and work on educating the sales staff.

Today, you get announcements and there's no way that our two marketing reps can tell me everything I need to know that's pertinent to my account and have the depth and the knowledge of what we need to answer questions about how we're going to implement ATM. ▄

"IF I TELL YOU HOW MUCH MONEY SYNC'S CENTRAL SITE SNA CONVERSION SAVED ME AT ████████, I'M DEAD MEAT."

Get the competitive tool companies would rather keep to themselves. The Central Site ConversionNode™ from Sync Research lets you downsize or eliminate FEPs now, and painlessly open your SNA network, with a platform for the further evolution of internetworking. It's available only from Sync, the most trusted name in SNA/LAN internetworking for the IBM customer.

"Once we saw the power of Sync, we told the reps from ████ and ████████ to get lost."

Short-term, the ConversionNodes offer potential savings of literally hundreds of thousands of dollars in reduced front-end processor costs. Longer term, by opening SNA, your huge SNA network investment can evolve gracefully, adopting key new technologies and standards like frame relay or APPN as they become appropriate.

"Our headquarters in ████████, was the hub of our SNA network. The trick was to increase access to our SNA hosts without killing ourselves financially. Adding ██ front-end processors to our network was going to cost ████ up front, plus ████ every six months in service."

SDLC requires significant FEP resources. But a single ConversionNode can replace up to 28 direct connections, supporting—via multidrop—hundreds of PUs. You can use much smaller FEPs or, in remote concentration applications, replace them entirely.

"The guys from Sync met with me and my boss at ████. They showed us how to save about $████ over the solutions from kingpins like ████, and the self-proclaimed ████ gurus."

Sync specializes in large-scale SNA connectivity. With over a decade of experience in IBM environments, Sync is expert in maximizing availability and minimizing costs.

"The others are just trying to slap your SNA sessions onto the LAN. Managing my ████ SNA traffic is too important to trust to that kind of a kluge. We just said, 'If we've got to do it, let's do it right.' Now our billing people in ████ can connect with ████ data located in ████ with the level of reliability we need."

"This technology is our secret weapon. Hey, you're not recording this, are you?"
Network Manager, ████ Inc.

ConversionNodes are 100% compatible with token ring, RS-232, V.35 and V.11 interfaces. They can be fully managed from the NetView console, are NetSpy compatible, and are as tightly integrated into SNA as native IBM devices.

"End result? This puts us head and shoulders above our key competitors: ████ and ████ I just hope they don't find out about it."

Get the free information kit "Short-Term Payback, Long-Term Thinking," for the whole story. Because you probably won't hear it from anyone else. Call us at 1-800-355-6355, Dept. E66.

SHORT-TERM PAYBACK LONG-TERM THINKING

SYNC RESEARCH
INTERNETWORKING SNA

Diagram: Host — FEP — Token Ring — Central Site ConversionNode™ — SDLC Lines

You say SNA routing

Router vendors add their own wrinkles to IP encapsulation, other schemes; and of course, there's always traditional SNA routing.

I say routing SNA

BY ANURA GURUGE

Multiprotocol bridge/routers are capable of many feats, but SNA routing is not one of them.

Contrary to vendor claims, this is an unequivocal fact regarding bridge/routers, whether they hail from Cisco Systems, Inc., Cross-Comm Corp., IBM, Proteon, Inc., Wellfleet Communications, Inc., 3Com Corp. or others.

A counter claim that Systems Network Architecture is not routable is equally false. SNA is an eminently routable protocol — provided, of course, that as with any other protocol, the appropriate routing software is used. SNA's equivalent of routers — IBM's 37XX communications controllers with Advanced Communications Function/Network Control Program (ACF/NCP) software — have been routing SNA traffic between 3270 terminals and applications resident in geographically dispersed IBM mainframes for the last 17 years.

What bridge/routers are capable of doing is routing SNA traffic, which is not the same as SNA routing (see story, page 58). This is not a question of mere semantics. Ignoring the distinction between these two terms could wreak havoc by dramatically altering the SNA traffic patterns in networks involving multiple physically dispersed mainframes.

An entire industry has been built on SNA routing, and yet another for routing SNA traffic. Nearly every router vendor today has its own twist on routing SNA traffic. With so many contenders, the issue is not as simple as, which one is better than the next? More precisely, the issue is which of the approaches is best tailored to a user's networking environment. And that is an issue worth exploring in greater detail.

LAYERED APPROACH

Routing SNA means bridge/routers must transport SNA traffic between an SNA local-area network gateway and an SNA device — between a LAN-attached personal computer emulating a 3270 device and a 37XX front-end processor — using a Layer 2 routing mechanism such as Internet Protocol encapsulation. SNA works above Layer 2, making it independent of and operating above the Data Link Control layer.

SNA routing, by contrast, is best thought of as occurring in Layer 4 and Layer 5 of the SNA stack. The transport mechanisms used by bridge/routers to route SNA are thus outside the scope of SNA and Layer 2 routing criteria that is not based on SNA. Hence, the end-to-end routes used in a multiple mainframe network by bridge/routers to transport SNA traffic may differ — sometimes significantly — from those that would have been used if SNA routing was used. SNA traffic will, in general, still reach intended destinations. But congestion may occur, in some instances, if traffic is routed over circuitous routes.

This type of problem is mitigated though, since only about 20% of the 50,000 or so worldwide SNA networks have multiple mainframes or remote 37XXs that use SNA routing. In networks with only one mainframe, or only channel-attached 37XXs, the inability of bridge/routers to emulate SNA routing is not an issue. In such networks, a single logical point-to-point connection is used between the remote SNA peripheral devices and a single data center containing the mainframe and 37XX.

Customers that currently rely extensively on SNA rout-

...ing should design their new networks cognizant that bridge/routers cannot emulate the routing that is done for them by remote 37XXs. If there is a continuing need for genuine SNA routing between mainframes, users fortunately still have some options. Remote 37XXs for SNA routing could still be used in conjunction with bridge/routers, with the bridge/router cloud providing a transparent Layer 2 network between the 37XXs. Advanced Peer-to-Peer Networking (APPN) and LAN-over-SNA routers could provide solutions. It is, however, also likely that some customers may decide to forego SNA routing and design their new network to route SNA based on a non-SNA scheme such as IP encapsulation.

APPN TO THE RESCUE?

APPN Network Node routing — now available on bridge/routers from IBM and 3Com — also does not equate to SNA routing, although it could be made to approximate SNA routing. APPN is based on session layer, SNA/APPN Logical Unit (LU) names. LU names are also the root route selection criteria of SNA routing. In multiple mainframe networks, APPN routing could be used, with appropriately predefined APPN class-of-service route selection tables, to approximate SNA routing. Unfortunately, APPN routing is not going to be the universal solution for customers requiring SNA routing — at least not for the time being.

In order for bridge/router-based APPN routing to take over the role of SNA routing, APPN has to be able to handle traditional SNA traffic — such as 3270 terminal-oriented data — rather than just program-to-program LU 6.2 traffic. The good news is that the so-called Dependent LU (DLU) technology required to do this will be available on mainframes, 3174s and PCs running OS/2 by September at the latest.

The bad news is that 3174s and OS/2 PCs only account for 50% of the current 3270 installed base. It is as yet unclear as to when DLU support will be available on non-IBM control units, gateways (such as Novell, Inc.'s popular NetWare for SAA), 3270/SNA emulators or, for

Who does what	
A guide to SNA routing providers	
Service type	**Supporting vendors**
SRB	Ubiquitous
IP encapsulation	
Proprietary	Cisco, 3Com
DLSw	IBM, Proteon, Wellfleet
PIR	CrossComm
APPN NN	IBM, 3Com
TDM	HyperCom Corp., Network Equipment Technologies, Inc. and others
Synchronous passsthrough	Cisco, CrossComm, Wellfleet
Remote SDLC polling	Cisco
Non-DLSw SDLC-LLC2 conversion	Cisco, CrossComm, NetLink, Sync Research, 3Com
AnyNet	IBM
RFC 1490 *	Eicon Technology Corp., HyperCom, NetLink, Sync Research
LAN-over-SNA	Adacom Network Routers, Computer Communications, Inc., IBM, Novell

*RFC 1490 is an Internet standard for routing SNA and bisynchronous data over frame relay circuits.

SOURCE: ANURA GURUGE, NEW IPSWICH, N.H.

Continued from page 53

that matter, on IBM's own Application System/400s and RISC System/6000s. It is unlikely to happen until at least 1995.

This could lead to an intriguing chicken-and-egg conundrum. The lack of third-party DLU support will preclude many SNA customers from being able to switch to bridge/router-based APPN. This perceived lack of customer endorsement for APPN could, as is the case now, make vendors reluctant to invest valuable resources implementing DLU on their products. The end result is that APPN routing as an alternative to SNA routing is still much more of a promise than a reality.

TECHNIQUES FOR ROUTING SNA

Bridge/routers have to contend with two distinct classes of SNA traffic: traffic to and from LAN-attached SNA devices such as PCs, 3174s, token ring-attached 37XXs, and traffic to and from link-attached SNA devices such as 3174s and 37XXs with serial ports. They use different techniques to cater to these two types of SNA traffic; this differentiation is another area of acute confusion. Some vendors, and many customers, think of bridge/router-related SNA routing as only applying to Synchronous Data Link Control link traffic. Not so. With LAN-attached PCs displacing 3270 terminals, and all of IBM's strategic products now being LAN-attachable, it is not unusual for at least half of the SNA traffic in typical IBM shops to be LAN-centric.

Today, there are three predominant techniques used by bridge/routers to route SNA traffic operating over LANs: bridging, IP encapsulation and CrossComm's Protocol Independent Routing (PIR).

IBM's Data Link Switching (DLSw) facility, which is now also available on IBM, Proteon and Wellfleet bridge/routers, is another option that does IP encapsulation of SNA, APPN and Network Basic I/O System traffic *among* other things. Encapsu-

lating SNA traffic directly in frame relay packets according to RFC 1490 is a more exotic technique just beginning to appear on the market. And SNA traffic, like any other type of traffic including voice and video, can obviously be transported end-to-end on a logical point-to-point basis using time-division multiplexing.

MAC ATTACK

Bridging, IP encapsulation and PIR all use the media access control (MAC) address of the destination SNA device as the basis of their SNA routing. Thus, SNA traffic is transported from MAC address to MAC address, regardless of SNA addresses or SNA application (or LU) names. The RFC 1490-based frame relay solutions, due to requirements of frame relay, do their routing based on preassigned frame relay virtual circuits, which will be assigned between the relevant SNA devices. Frame relay may also be used between bridge/routers, independent of RFC 1490, to transport traffic — including SNA traffic. In such situations, the bridge/routers would use source route bridging (SRB), IP encapsulation or PIR across frame relay, typically using frame relay encapsulation schemes developed prior to RFC 1490.

SNA link traffic is currently supported by bridge/routers using three disparate schemes. These being: synchronous passthrough, remote polling (via such techniques as local SDLC termination or poll spoofing) and SDLC-to-Logical Link Control 2 conversion. Synchronous passthrough and remote polling both work on a serial link-to-serial link basis — for example, an SDLC link-attached 3174 connected via an internetwork to a serial port on a 37XX. Both schemes work by having some form of predefined, port-to-port, point-to-point mapping between the two bridge/router ports — one at the host site and the other at the remote site — to which the two links are attached.

Once this mapping has been established, which is sometimes realized by assigning pseudo MAC addresses to the serial ports, it is possible to use IP encapsulation, PIR, a simple High-Level Data Link Control encapsulation or frame relay encapsulation to transport the SDLC data across the internetwork.

SDLC-to-LLC2 conversion, however, is now rapidly becoming the preferred approach for supporting SDLC link traffic. With this conversion, which can now be done directly within bridge/routers without recourse to external conversion boxes such as those from Sync Research Corp. or NetLink, Inc., SDLC devices are made to appear as if they are LAN-attached devices. Thereafter, the same routing techniques used for LAN-attached SNA devices — bridging and IP encapsulation — can be used to deal with traffic to and from the SDLC-attached devices.

PERIPHERAL TECHNIQUES

IBM's AnyNet technology, now available on mainframes, AS/400 and OS/2 machines, is another nascent approach for routing SNA traffic across a LAN/WAN internetwork. AnyNet is a protocol conversion methodology that converts SNA traffic to bona fide Transmission Control Protocol/Internet Protocol traffic replete with TCP/IP addresses. Bridge/routers will shuttle the TCP/IP traffic from the source AnyNet converter to the destination AnyNet converter. The destination AnyNet device will then convert TCP/IP data back into SNA.

Though they do not have to worry about routing SNA per se, LAN-over-SNA routers also need to be considered in relation to the issue of SNA routing. LAN-over-SNA routers are, in effect, the antithesis of the IP encapsulation approach. These routers encapsulate IP, Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) and NETBIOS traffic within SNA LU 6.2 frames and transport the LAN traffic from one LAN to

Comparison of SNA routing techniques

	SNA routing	APPN NN routing	SRB	IP encapsulation	Current DLSw	CrossComm's PIR	IBM's AnyNet	LAN-over-SNA
Routing criteria	SNA application/ LU names	APPN application/ LU names	MAC address	MAC address	MAC address	MAC address	APPN application/ LU names	IP and IPX addresses; NETBIOS names
SNA routing	✔	(1)						✔
Degree of openness	De facto standard	IBM license required	De facto standard	Vendor-specific	De facto standard (2)	Proprietary		Vendor-specific
Fixed path	✔	✔	✔					✔
Destination location	Predefined or broadcast search	Predefined or broadcast search	Broadcast search	Predefined, broadcast search or learned	Broadcast search	Broadcast search or learned	Predefined or broadcast search	Broadcast search, predefined or learned
Destination location caching	✔	✔		✔	✔	✔	✔	✔
Routing technique	Subarea-to-subarea route selected from predefined tables based on a single-value class of service associated with each session.	Route calculated using network topology database using a class of service that can consist of up to 9 criteria. A spanning tree could also be used.	Route is taken by the first response to the search selected. This route, in the form of a RIF, is prefixed to all message units.	IP address of the bridge/router adjacent to destination is used to route TCP/IP datagrams using standard IP routing.	IP address of the bridge/router adjacent to destination is used to route TCP/IP datagrams using standard IP routing.	Dynamic routing scheme, not that dissimilar to IP, that takes into account route congestion and failures.	IP routing across the network.	SNA address of the bridge/router adjacent to destination used to route LU 6.2 frames containing LAN traffic using standard SNA/APPN peer-to-peer routing.
Strengths	• SNA's native routing scheme. • Based on session-layer names. • Used in around 10,000 networks. • Transmission priority. • Load balancing. • Parallel links for protection against link failure. • Congestion control.	• Plug-and-play. • Sophisticated, multiple criteria-based route selection. • Transmission priority. • Congestion control.	• Plug-and-play. • Proven de facto standard.	• Dynamic alternate routing. • Load balancing. • Transmission priority. • No repetitive searches. • Potential for criteria-based path selection.	• Dynamic alternate routing. • Load balancing. • No repetitive searches. • Potential for criteria-based path selection. • Proteon implementation close to plug-and-play. • Includes SDLC-to-LLC2 conversion and local LLC2 acknowledgments	• Plug-and-play. • Dynamic alternate routing. • Load balancing. • Transmission priority. • No repetitive searches. • PIR II supports network partitioning.	• Dynamic aternate routing. • Load balancing. • Transmission priority. • No repetitive searches. • Potential for criteria-based path selection.	• SNA-based LAN interconnection scheme. • Leverages all of SNA's features, including reliability and congestion control.
Weaknesses	• No dynamic alternate routing. • Convoluted and complex.	• No dynamic alternate routing. • No parallel links. • No load balancing. • Does not adapt to new routes that become available.	• No dynamic alternate routing. • No parallel links. • No load balancing. • Repetitive searches. • Does not adapt to new routes that become available.	• Not plug-and-play. • Vendor-specific.	• Standard version not plug-and-play. • No traffic prioritization.	• CrossComm-specific.	• Essentially an IBM-specific prototype. • Limited to 3 platforms. • Uses mainframe cycles for protocol conversion.	• Only supports IP, IPX and NETBIOS. • Without frame relay, performance could be an issue. • No market leaders.

Footnotes:
(1) Can be approximated.   (2) Multivendor standard in development.   ✔ = Supported

SOURCE: ANURA GURUGE, NEW IPSWICH, N.H.

Imagine you had the same choice in internetworking partners.

Which one would you select?

At Wellfleet, full service means you get more. First, a partner who listens to your needs and architects the right internetworking solution for you. Second, our family of multiprotocol routers and multiservice WAN switches offers the industry's highest availability and performance with support for every network environment from SNA to ATM. And finally, Wellfleet provides worldwide support for your internetwork every hour of every day. So get the service you fully deserve. Give Wellfleet a call at 1-800-989-1214, extension 32.

WELLFLEET

The new AS/400.
Black goes with everything.

If your business plan includes client/server computing, rest assured that the new AS/400® Advanced Series will coordinate impeccably. Everything about it is designed with a new openness and interoperability, to thrive in client/server environments. There's open client access, to serve OS/2,® Windows,™ DOS, Macintosh or UNIX® systems with equal ease and security. Powerful features that provide open database access and protocol-independent networking are built in. There's a wide choice of AS/400 Advanced Servers—high performance processors fine-tuned for client/server duty. And thousands of AS/400 client/server applications are ready today. In short, the new AS/400 is so easy to implement and manage, it represents a new generation in client/server computing: a system designed to stand out, even as it blends in. For additional information, call 1 800 IBM-6676, ext. 641.

Processor Activity

1 Low High
2 Low High
3
4

The new AS/400 Advanced Series from **IBM**

another across an SNA network.

LAN-over-SNA routers are bona fide APPN Type 2.1 peer-to-peer nodes that are directly connected to an SNA network and act as transparent, physical unit (PU) passthrough gateways in relation to downstream SNA devices. Consequently, genuine SNA routing is possible in LAN interconnection networks built around LAN-over-SNA routers. The SNA routing, however, will be performed by 37XXs and mainframes, rather than by the routers themselves.

TEETERING ACROSS BRIDGES

The bridging of SNA traffic — which, in most cases, still tends to be token-ring's SRB though Ethernet-oriented transparent bridging may also be used — is supported by most bridge/router vendors. SRB relies on a dynamic broadcast search initiated by the source device to determine the location of a remote destination.

The search may be conducted on a single-route or all-route basis. SNA devices in general opt for all-route searches. The route traversed between the source and the destination devices is recorded within the Routing Information Field (RIF) of the explorer packets received by the destination; it is then returned to the originator via search responses issued by the destination. The originator typically decides to use the route indicated by the first response it receives and inserts the RIF contained in

Will the real DLSw please stand up?

The status of the so-called Data Link Switching (DLSw) standard is in some disarray. DLSw implementations available today are based on the Internet's RFC 1434, which reflects IBM's original implementation of DLSw on the vendor's 6611 router.

There is, however, a multivendor group working on an enhanced DLSw standard that will be ratified by the Internet Engineering Task Force (IETF) when it is complete. This new DLSw is supposed to offer improvements in the areas such as flow control, traffic prioritization, standard management and load balancing.

Cisco Systems, Inc. and 3Com Corp., whose current Internet Protocol-encapsulation schemes already address many of these features, are understandably reluctant to implement the current DLSw, which they view as being regressive.

They intend to wait for the IETF version of DLSw, but this may not be available until late this year. Proteon, Inc.'s enhancement to the current DLSw creates yet another subspecies of DLSw to add to the confusion. And if this isn't bad enough, the powerful congestion control mechanism of DLSw, which is based on Advanced Peer-to-Peer Networking's Adaptive Pacing, happens to be a part of the patented APPN technology. Whether DLSw implementations will thus require an APPN-type technology license to use adaptive pacing has yet to be conclusively resolved.

Sometime in 1995, DLSw will emerge as a key multivendor interoperability standard for SNA internetworking. Unfortunately, there are quite a few hurdles to be negotiated prior to that.

Benefits of DLSw

Data Link Switching (DLSw) provides:

IP encapsulation of SNA, APPN and NETBIOS.

SDLC-to-LLC2 conversion.

Local LLC2 acknowledgment that precludes the LLC2 Receive Ready responses that are used to acknowledge every SNA frame transmitted across WANs.

MAC address and NETBIOS name caching.

SNA routing defined

SNA routing is Systems Network Architecture's native routing scheme. It is a fixed-path routing scheme based on manually predefined routes.

SNA routing only occurs between 37XXs, between 37XXs and mainframes, or between mainframes; it is not supported between peripheral devices such as a 3174 control unit and a 37XX.

The technology behind SNA routing is based on SNA subareas and SNA virtual routes between subareas. An SNA subarea can only be formed around an SNA Type 5 or Type 4 subarea node. Advanced Communications Function/Network Control Program (ACF/NCP) in a mainframe implements a Type 5 node, while ACF/NCP in a 37XX implements a Type 4 node. For a bridge/router to be able to perform SNA routing, it would have to contain the software that would implement a

Type 4 node.

To its credit, Cisco Systems, Inc. recognized the role of SNA routing and explored the possibility of adopting the technology back in 1991 as Phase V of its original five-phase SNA internetworking strategy. There is a significant amount of real-time processing, buffer space and memory for software code required to realize a Type 4 node. In addition, some of the newer protocols used by subarea nodes are not in the public domain — since IBM started in 1985 to progressively shift SNA from an open architecture at the time to a more closed and proprietary approach. Thus, implementing a Type 4 node on a Cisco bridge/router, particularly during 1991 and 1992 when Cisco was still saddled with a single-processor scheme, was far from easy. Recognizing the difficulties, Cisco abandoned this initiative in early 1992.

---

that response in all subsequent frames transmitted.

SRB, which is a standard feature of the token-ring architecture, is a universally applicable technique for transporting traffic between LANs, independent of their Layer 3 (and above) protocols.

SRB unfortunately is riddled with frustrating flaws that severely depreciate its appeal and applicability. For starters, it is a fixed path routing scheme that does not provide for dynamic rerouting of traffic across alternate paths in the event of a path failure. Neither does it support load balancing across multiple paths. Based on the first response received, the path chosen by SRB may, in some cases, not be the optimum path between the source and the destination. There is no criteria-based mechanism for path selection, such as selecting the path with the least number of intermediate hops.

In addition, SRB does not cache the routes to previously discovered destinations. Thus, repeated broadcast searches, which in large nets could generate a significant amount of spurious overhead traffic, could be continually performed for the same destination. These repetitive searches are particularly galling in SNA environments where much of the traffic is between LAN-attached PCs and a single 37XX-type LAN gateway to the mainframe. A standard SRB bridge on a LAN with 30 PCs, for example, will unashamedly do 30 broadcast searches for the LAN gateway. And this repetitive search for the mainframe LAN gateway will, in turn, be done by every bridge on the net.

BURNING BRIDGES

IP encapsulation and PIR attempt to comprehensively eradicate the shortcomings of SRB. They both set out to provide dynamic alternate routing and load balancing, eliminate repetitive searches and offer a level of criteria-based path selection. Nonetheless, both techniques — just like SRB — rely on the Layer 2 MAC address of the destination SNA device as their basic route selection metric. The differences between these two techniques, as well as the various flavors of IP encapsulation, stem from the mechanisms used to locate the destination MAC addresses.

With IP encapsulation, each destination MAC address has to be associated with an IP address for the bridge/router connected to the LAN on which the destination device resides. Bridge/routers pass IP datagrams, which contain SNA message units, between the source LAN and the destination LAN using the IP address of the destination bridge/router. There are two distinct approaches that can be used to correlate destination MAC addresses with IP addresses. A user could manually define at each bridge/router a list of destination MAC addresses and the corresponding IP addresses. This is the approach currently used by 3Com. If there are only a few SNA destinations — a single mainframe LAN gateway, for instance — this is a simple and acceptable approach. But faced with a slew of destination devices, manually setting up the route definitions is an ungainly, if not downright unacceptable, approach.

Cisco's routing software, IBM's DLSw and CrossComm's PIR all use a dynamic automatic search mechanism, similar in overall concept but not in execution to that used by SRB and APPN, to locate destination MAC addresses. Again, the amount of manual predefinition required by these techniques varies significantly. PIR, which is not IP-based and therefore does not use

---

IP addresses, is essentially a plug-and-play scheme that does not require the additional definitions required by SRB bridges. This results in less work setting up addressing tables.

The Cisco scheme and IBM's standard DLSw require that the IP addresses of the other partner bridge/routers are defined at each source bridge/router. This results in more manual intervention than PIR, or even Proteon's approach.

Routing SNA across subareas

In hybrid LAN/SNA internets, bridge/routers replace front-end processors (FEP) as the pivotal devices for forwarding SNA and LAN traffic. Since these devices do not perform SNA routing, traffic will reach end points, but some routes may be congested while others are hardly taxed.

GRAPHIC BY SUSAN SLATER   SOURCE: ANURA GURUGE, NEW IPSWICH, N.H.

Proteon, in its implementation of DLSw, has eliminated this requirement by using DLSw-specific IP multicast addresses to locate other DLSw routers. This enhancement makes DLSw pretty close to plug-and-play and, consequently, dilutes PIR's hitherto significant edge when it comes to simplicity and administrator friendliness.
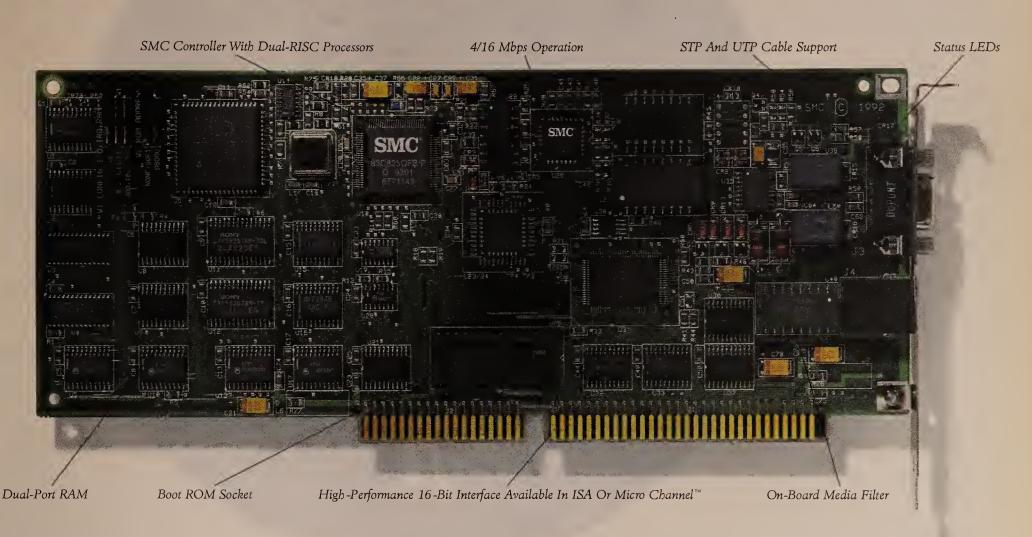
ROUTE'S END

Bridge/routers, though capable of adroitly transporting SNA traffic between LANs using non-SNA Layer 2 techniques, cannot perform genuine session-layer SNA routing as performed by IBM 37XXs or mainframes. Fortunately, this is not an issue for the 80% or so of the SNA customer base that does not rely on SNA routing. It can, however, be a major impediment to those that do and were planning to migrate to a bridge/router-based LAN/WAN internetwork.

One solution for such customers is to use remote 37XX for SNA routing in conjunction with a bridge/router net. APPN and LAN-over-SNA routers are also potential options. The lack of across-the-board support for 3270 devices in APPN does diminish its applicability. It may, in some cases, also be possible to design new bridge/router-based networks that do not require SNA routing. The bottom line in all cases is that when designing an SNA-capable bridge/router network, users have to always be cognizant that routing SNA is not the same as SNA routing.

➥ Guruge is an independent consultant specializing in internetworking and IBM network architectures. He writes extensively, presents seminars worldwide and can be reached at (603) 878-1303 or via Internet/MCI Mail at aguruge@mcimail.com.

A

They cost the same. The
You call the 800 number at com
Hmmm. What to

The call goes to Company B, of course.
As well as the sale. And the next sale, too.
Moral of the story? When the product you're
selling is parity, make sure your service isn't.

And one way to do that is with a more
comprehensive, better planned communications
system. At Ameritech, we have a long history
of creating solutions for a variety of business

B
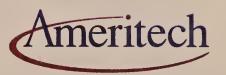
y meet the same specs.
pany A and get a busy signal.
do? What to do?

problems. And with our broad range of communications products and services, we can design a solution that will keep your company accessible to

all of your customers all of the time.

Welcome to the passing lane on the information highway. Just call 1-800-719-5822, extension 15.

Ameritech

Your Best Link
To Better Communication

Letters

Security check

I adamantly agree with Mark Gibbs' point that hackers don't use magic tricks to enter a computer system. Any entry point into a system via the Internet, dial-up modems or an electronic data interchange connection is susceptible to entrance by an intruder.

Passwords are never guaranteed, but they do provide a first line of defense if proper security management is utilized.

A second line of defense is auditing. An organization should review its logon reports regularly to determine who is attempting to access the system, noting failed attempts by ID and access point. Surprisingly, although hackers are a threat, most computer crimes are committed by disgruntled employees.

With the growth of today's distributed, interconnected networks, passwords may become compromised. Other security measures need to be taken to protect specific entry points and supplement the security program. User authentication, in which users are verified by security devices like a hand-held token, smart card or biometric device (such as a signature or thumbprint), guards the ever-expanding elements of a company's information system.

As we connect to the Information Superhighway, we need to consider that security will only be as strong as the weakest link. The more entry points into a system, the more vulnerable it becomes. Corporate America should not resign itself to the fact that "hackers will be hackers," but work constantly to improve information security before unauthorized access occurs.

Ralph Massaro
General manager
CKS North America
Pittsburgh

Hackers vs. crackers

Regarding Mark Gibbs' column "Heavens! Hackers, horror and the Highway" (May 16, page 24):

Gibbs should obtain a copy of either the *Hacker's Dictionary* or the Jargon file (which I know is on-line both on the Internet and on CompuServe). He consistently used the term hackers in his article to refer to those who crack into other's computer systems without permission. He should have referred to these people as crackers.

I know Gibbs stated that he was using the term in the perjorative sense. But if you check the Jargon file, you will find that this perjorative sense carries a disclaimer — that this meaning of the word is only used by those who do not know any better than to misuse the language.

Given that the Internet itself is one beaut of a hack, please don't knock those of us out here who thoroughly enjoy using our skills to further the fields of information technology, programming and networking by creating "hacks," which are then used by hundreds, if not millions, of others.

Yes, I'm a hacker (under the official definition) and proud of it. I'm also a network administrator and security officer for the U.S. government. I do not break into systems for fun or profit. In fact, some of my better hacks are those which prevent crackers from entering my systems, including my home bulletin board system.

Steven Foust
Microcomputer specialist
Department of Agriculture
Glen Ellyn, Ill.

Help desk

IBM announced its Turboways 100ATM Adapter in April. It costs $1,995. For more information, call (800) 426-2255.

HP and Silicon Graphics have yet to detail ATM adapter card plans.

Is there a way to get a network running Windows, Windows for Workgroups and Windows NT to talk to a VAX/VMS system? I need to access files on the PC from the VAX and vice versa. All of these systems are on a single Ethernet. The goal is to gradually migrate everything to the Windows NT box, so we obviously don't want to spend too much.
Richard Fisher, Lubbock, Texas

Brad Cooper, consultant at Metamor Technologies, Ltd., a software consultancy and systems integrator in Chicago, replies:

A seamless way to access files on your VAX from your Windows for Workgroups and Windows NT nodes would be to install Pathworks Version 5 for Open VMS (LAN Manager) and use NETBIOS Extended User Interface from your Windows for Workgroups and Windows NT clients. You will need a Pathworks VMS TK50 Media and Documentation Kit (part number QA-A93AA-H5) at $249 and a Pathworks (LM FPS VMS SVR) File and Print License (part number QM-A93AA-AB for 10 concurrent clients) for $3,000. These items can be ordered directly from Digital by calling (800) 344-4825.

Accessing the PC files from the VAX is a little trickier, involves installing additional components on your personal computers and probably takes more time than you are willing to invest. Some options that we use include: DECnet and its Files Access Listener utility (Pathworks Combination Client Server license) and Vector Networks LANutil product, which is available from Banner Computers, Inc. at (617)938-1400.

Threat from within

I read with some amusement Mark Gibbs' column about hackers. What is especially funny about the CalREN episode is that they're worried about their network being penetrated from the outside, while giving access to thousands (hundreds of thousands? millions?) of high school students. They will soon discover that they need world-class security, anyway.

Dale Worley
President
NetGuide Information Services
Waltham, Mass.

Slow down

You published a Buyer's Guide on high-speed modems (April 4, page 93). Have you done a similar review of low-speed modems? If so, please give the issue date. Thank you.

Steve Des Palmes
Manager of telecommunications
Ball Corp.
Boulder, Colo.

Editor's reply: The chart accompanying the Buyer's Guide to high-speed modems that ran in our April 19, 1993 issue (page 44) noted modems that support speeds of 1,200 bit/sec to 28.8K bit/sec.

Can you get to the next generation of smart networking without scrapping the old one?

Right now, lots of vendors are selling next-generation smart switching hubs. But they all have one minor problem: What they sold you before won't work with their new technology.

With Chipcom, it's a different story.

Because to us, all this new technology is technology we pioneered years ago. We've been building switching hubs as long as we've been building hubs. So if your network has been built around Chipcom products, you've already bought a big chunk of its future. And whatever technology that future will require.

ONdemand™ NCS

For example, if you're running a network on our ONline™ System Concentrators, and you need to upgrade all or part of it to our next-generation ONcore™ Switching System, there's no problem. Every ONline module is completely upward-compatible. So not only do you save substantially by not scrapping your existing modules, you can upgrade where and as needed.

And unlike hub makers who grew up struggling with departmental LANs, vast enterprise-wide networks don't faze us. We started out building connectivity devices for huge networks, so we built scaleability and

With Chipcom® you can.
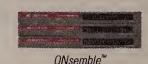
comprehensive network management into our architecture right from the start.

Call us at **1-800-228-9930** for complete information on the ONcore Switching System.
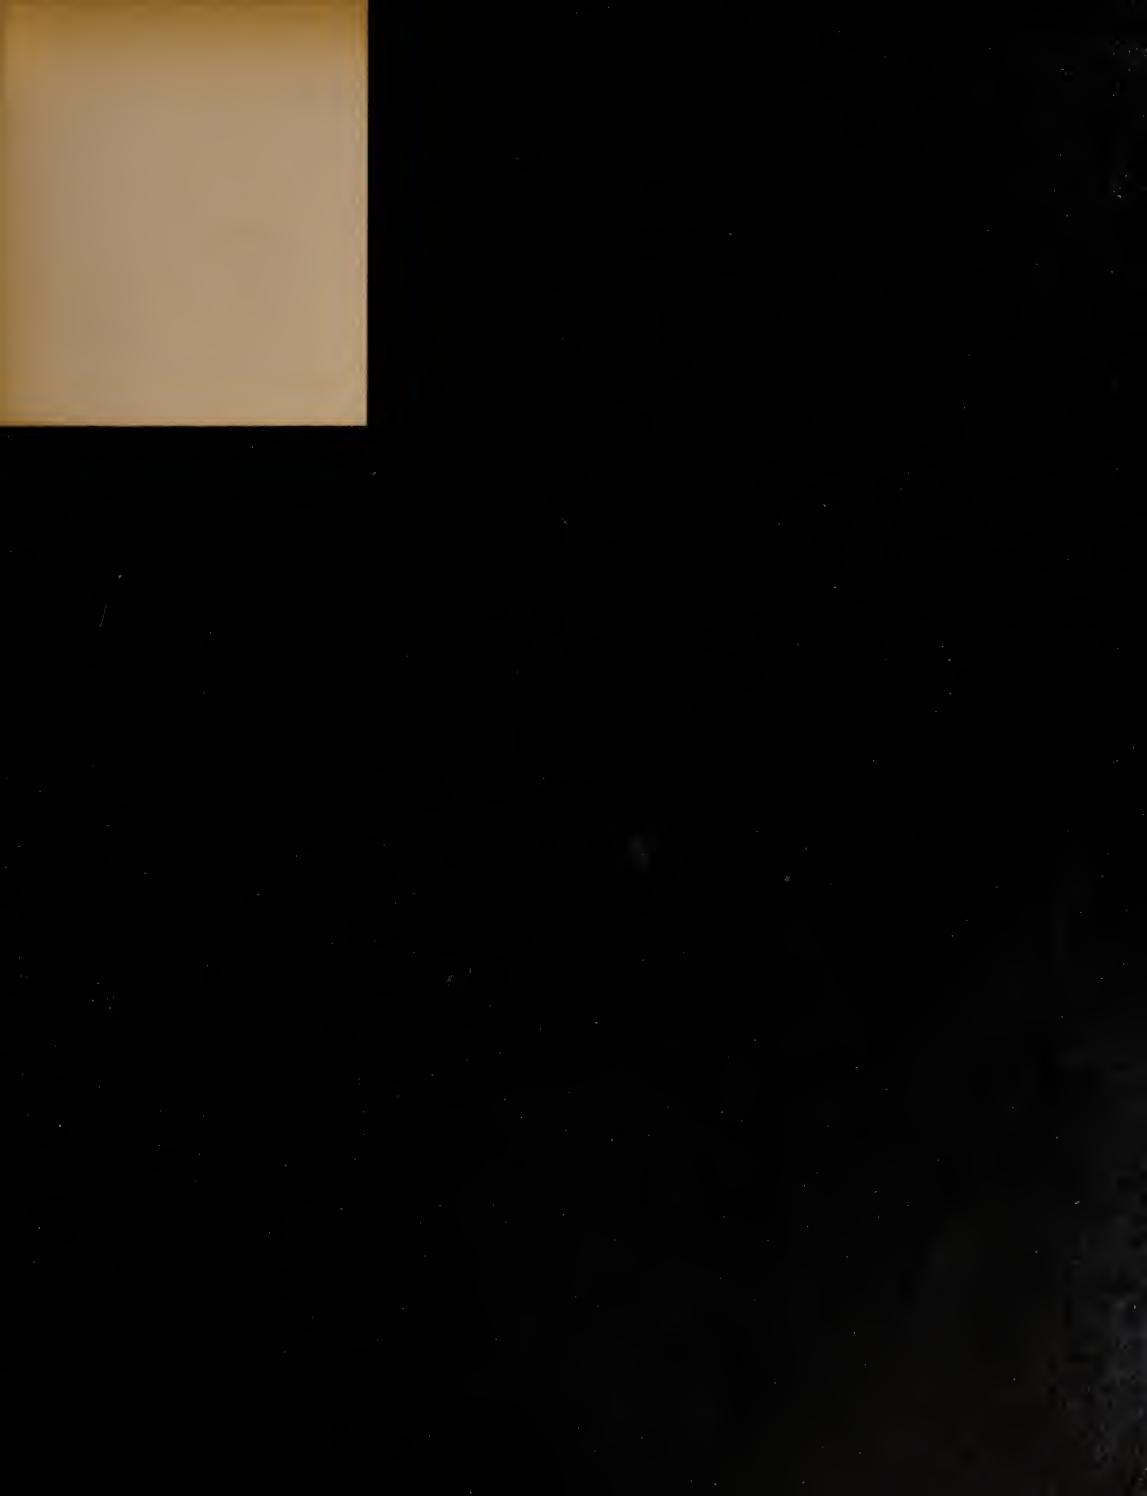
Or call us at the same number to inquire about Chipcom's worldwide seminar series on enterprise-wide networking. It's being held in more than 50 locations, so call to learn where and when we'll be in your region.

ONcore™   *ONline™*   *ONsemble™*

◆ **CHIPCOM**™

APPN evolves
to a

Advent of High Performance Routing should pit APPN more evenly against TCP/IP in tussle for backbone dominance.

higher order

BY THOMAS ROUTT

T he differences between IBM's Advanced Peer-to-Peer Networking and the Transmission Control Protocol/Internet Protocol suites are like night and day to some people.

IBM's evolution of APPN, however, appears destined to obscure some of the differences between the two and draw APPN closer than ever in functionality to TCP/IP. IBM, in fact, is transforming APPN into a hybrid architecture, which, when implemented, will be capable of supporting connection and connectionless services.

Currently, APPN is a connection-oriented service only, an aspect of the architecture that may carry some performance penalties for router-based internet users. That's because connection-oriented services reserve resources at hops across a network to ensure an end-to-end connection. But with the availability of High Performance Routing (HPR), APPN will gain the intelligence to mitigate and even stem congestion on network backbones, a feat that would enable APPN to gain an advantage over TCP/IP.

Systems Network Architecture and TCP/IP trace their development from quite separate origins and philosophical bases, yet today, they each support distributed, client/server applications as well as host-centric connections. Both have been designed to provide local as well as wide area connectivity.

Philosophical differences aside, the decision to use APPN vs. TCP/IP is not as clear-cut as night and day. Decision points weigh heavily on users' existing network environments, traffic types and cost issues, among other things.

Perhaps one of the more intriguing areas of comparison given recent technology developments is that of routing and connection/connectionless service philosophies. But before delving into any analysis of APPN and TCP/IP routing infrastructures, it pays to understand how IBM's technology has developed to its current stage.

SNA'S ROOTS

SNA was introduced by IBM in 1974 as a host-centric, hierarchical (subarea) network and has evolved over time, taking on peer-to-peer characteristics along the way (see graphic).

During 1985, IBM introduced a Low Entry Networking (LEN) node, providing peer-to-peer functionality in point-to-point links supporting LU 6.2 interprogram sessions. APPN/Intermediate Session Routing (ISR) was prototyped on the System/36 during 1986, and IBM introduced a network node control point (NNCP) and end node control point. The NNCP provides several server functions, including dynamic topology exchanges, directory updates and optimum path calculation.

Today, APPN/ISR is supported on all of IBM's major platforms, including System/390 and System/370 hosts, 3745s (the host and 3745 APPN support is aggregately referred to as a composite network node), Application System/400s, RISC System/6000s, Personal System/2s (OS/2 and DOS), and 3174s. More than 40 non-IBM vendors have also announced APPN products or directions.

The next step in APPN's evolution is HPR, which represents a significant technology improvement over ISR products. IBM is expected to deliver HPR support beginning in the fourth quarter.

APPN/ISR ROUTING

APPN/ISR provides connection-oriented internodal links to support end-to-end logical connections between logical units in session. The ISR connection orientation at the link level between each pair of adjacent nodes is provided through a routing technique called label swapping.

Additionally, APPN/ISR provides a connection orientation not only for end-to-end logical unit-to-logical unit sessions, but also at the link level between adjacent nodes. This gets to the heart of the philosophical difference between APPN's connection orientation today. APPN fans find security in the connection orientation because it guarantees class of service levels and manageability. TCP/IP proponents, however, argue that the connection nature is inherently wasteful because it consumes resources that could be applied elsewhere.

ISR provides local form session identifier label swapping as well as session connector-based transport functions at each intermediate node along the session path. This approach is in line with an APPN major design point to provide reliable support for enterprise mission-critical applications over links that may range in quality from good to poor on a node-by-node basis.

The major disadvantage of APPN/ISR is that the loss of a link or node along the path results in loss of logical unit-to-logical unit sessions, requiring that a network node recalculate the path in order to support session reestablishment.

ENTER HPR

When it ships this year, APPN/HPR will augment APPN/ISR; it is designed to take advantage of fast link speeds with low bit error rates. The major benefits of HPR, also referred to as APPN + during its initial disclosure in March 1992, is its able to nondisruptively reroute sessions around failed nodes and links. HPR also will introduce rapid transport protocol (RTP) and automatic network routing (ANR).

RTP is a connection-oriented, end-to-end transport protocol that runs over ANR, the latter providing deterministic connectionless

TCP/IP is tops

TCP/IP is the protocol users rated as the most important for the next two years, while APPN ranked second, according to an International Data Corp. study.

SNA/APPN evolution

1974

SNA subarea

- Host-centric networks
- Program-to-device
- Program-to-program
- Hierarchical

LEN node

- Point-to-point
- Peer-to-peer LU 6.2 sessions
- Composite VTAM-NCP LEN node

1985-1987

APPN/HPR

- Peer-to-peer packet routing
- High-performance routing
- End-to-end flow control
- Error recovery
- Standard interfaces

1986-1993

Broadband network services

- High bandwidth
- Multigigabit links
- Multiprotocol routing
- Isochronous traffic
- Multimedia
- Bandwidth management
- Standard interfaces

1994-1995+

SOURCE: VEDACOM CORP., SEATTLE
GRAPHIC BY SUSAN SLATER

source routing over minimal data link control. In essence, RTP functions as a virtual link.

RTP provides a range of benefits, including connection awareness of each session running over the link; reordering, if necessary, following a path switch; reordering if the route contains one or more multilink transmission groups; optional reliable delivery; flow control and congestion control/avoidance; and nondisruptive route switching.

matic route calculation in APPN is one of the leading selling points for a migration from subarea SNA, which requires thorough, preconfigured path definitions.

But on a more microscopic level, the differences between APPN/HPR and TCP/IP begin to surface. Architecturally, APPN and TCP/IP can effectively implement any of three types of routing decision processes: destination routing, source routing and label swapping, which is also referred to as call mapping (see graphic).

Destination routing, source routing and label swapping		Originating node	Intermediate node
Destination routing	At session startup	Nothing	Nothing
Destination routing	For each packet	Examine packet header. Note destination address. Check table for next hop on best path to destination. Send packet to that hop.	Examine packet header. Note destination address. Check table for next link on best path to destination. Send packet to that hop.
Source routing	At session startup	Check table for best total path. Put total path information in packet header. Send packet on first hop.	Nothing
Source routing	For each packet	Add path information to each packet header.	Examine packet header. Note self in path information. Send packet to next hop listed after self.
Label-swapping	At session startup	(If the source determines route) Check table for best total path. Assign code for session. Put total path information in packet header. Assign label for first hop. Send setup packet through path.	Examine packet header. Note code for session (if source selects route) or destination address (if path not preassigned). Note label assigned for hop to self from previous node. Assign label for next hop. Enter previous-hop label/next-hop-label pair in table. Swap labels and send to next hop.
Label-swapping	For each packet	Check session code. Check table for label. Put label for first hop in header. Send packet on first hop.	Examine packet header. Note label on packet. Look up label pair in table. Swap prior label with next label. Send packet to next hop.

ANR runs logically under RTP and provides source routing with locally specified labels; discards incoming packets during congestion episodes; offers connectionless, stateless fast routing; and handles priority-based servicing of outbound transmission links.

HPR will use adaptive rate-based congestion control to compensate for the loss of the adaptive pacing function found in APPN/ISR intermediate routing nodes. Adaptive rate-based congestion control regulates the input traffic to the RTP logical link and is preventive, not reactive. Currently, both APPN and TCP/IP support reactive window flow-control schemes; users adjust window sizes after congestion is detected. By contrast, as a network approaches a congestion condition — as determined by increasing delay and diminishing throughput thresholds — adaptive rate-based congestion control reduces the input traffic rate until network capacity returns to acceptable levels.

Adaptive rate-based congestion control employs a closed-loop feedback mechanism based on information periodically exchanged between RTP components at the end points of an RTP logical link. The two feedback rates to watch are: the rate at which RTP accepts data arriving from the network and the rate at which RTP conveys data to a recipient. Monitoring of these two rates provides the basis for congestion anticipation and enables preventive throttling.

IP VS. APPN

The enhancements imminent in APPN/HPR will boost the protocol suite so it can effectively compete on a more even keel with TCP/IP, which has had a distinct advantage in terms of supporting connectionless class of service. Both TCP/IP and APPN are designed to perform automatic route calculation from topology information. The availability of auto-

Destination routing is used at the network layer by protocols such as IP and at the data link layer by Ethernet. Source routing is used at the network layer by protocols such as APPN/HPR and at the data link layer by token ring. Label swapping is used at the network layer by APPN/ISR and at the data link layer by Asynchronous Transfer Mode (ATM).

In destination routing, the originator sends out each packet with the destination address in the header. Each intermediate system reads the destination address, calculates the most efficient next hop and then forwards it. Destination routing advocates argue that each packet can find its best path at each intermediate node and end systems should not waste their resources on routing calculations.

In source routing, the originator discovers all routes to the destination and selects the optimal route. It then places a list of each intermediate node and the links between them in each packet's header, and each intermediate system notes the next link in the header and forwards the packet over it. Source routing proponents posit that source routing consumes fewer processing resources at intermediate nodes

because the path is already calculated at the initial node. Each intermediate node has to only note the next link listed in the header.

In label swapping, the route is calculated at the beginning of a session either by the originator or at each intermediate node. An initial packet is sent once over that path to let each intermediate node allocate its own identification for the next link for that session, and each successive packet arriving on one link for that session — as identified by the label from the previous node — is forwarded over the next hop with a new label for that hop. Label swapping supporters agree with source routing proponents that it is best to use a predetermined path. Label swapping also uses the shortest packet headers, decreasing net traffic. However, TCP/IP headers can be, and often are, compressed.

CONNECTION ORIENTATION

Those engaged in source vs. destination routing polemics often sidestep into issues not directly tied to the routing technique to be deployed. This leads to the routing technique itself being used inappropriately for decision making.

Source routing is usually used by connection-oriented protocols, while destination routing is generally used by connectionless protocols. But that is not always the case. APPN/HPR, for example, uses a connectionless routing protocol with a source routing process.

Each of the three types of routing can be implemented well or poorly. For example, APPN/ISR, which uses label swapping, performs more poorly than APPN/HPR. ATM, which is a very efficient protocol, also uses label swapping.

Both TCP/IP and APPN/ISR have connection-oriented services between pairs of communicating end systems. Of the two, TCP/IP offers connectionless service within intermediate systems, the benefit being it provides dynamic rerouting in the event of a failure. Currently, APPN/ISR only provides connection-oriented service, both in the end systems and in the hop-by-hop flows between intermediate systems.

A connection is defined as a data flow between communicating partners, which assures certain characteristics such as flow control, reliable delivery, sequential data delivery and no duplicate data delivery.

In a backbone, the core routing service differs between IP and APPN (see graphic). IP routing is connectionless; APPN/ISR routing is connection-oriented. APPN/HPR routing will be connectionless.

Consider the different connection characteristics of TCP/IP and APPN/HPR as backbone protocols. In an IP backbone, TCP connections are seen only by the end systems. The end systems inject datagrams into the network, while intermediate systems forward the datagrams. Outside of datagram forwarding, there is no requirement that determines delivery order for packets or defines reliability of packet delivery. Only the end-system TCP stacks handle connection requirements.

Currently, APPN/ISR is replete with connection-oriented facilities. For instance, point-to-point links between end nodes and network nodes, as well as network nodes to network nodes, are connection-oriented.

End node-to-end node sessions — such as LU 6.2 and LU 2 — are connection-oriented. No checksums or retransmission logic exists at the data link layer. For purposes of flow control, each end node-to-network node and network to network node hop is connection oriented. (IBM calls these hop-by-hop links session stages.)

A broader definition of connection might be argued. Under this broader definition, any environment that contains source routing or any reservation of resources qualifies as a connection-oriented system. APPN/HPR ANR forwarding would be considered a connection-oriented scheme. However, using this broader definition, IP would become a connection-oriented scheme since it has a source routing option. Isochronous extensions to IP would also be considered connection-oriented.

HPR is similar to TCP/IP — albeit with an important distinction. The backbone is connectionless, supporting source-routed datagrams with four priority levels — compared to three priority levels in an IP environment. The type of service — or class of service — is really not seen within the backbone, although the route server computes ANR routes with type-of-service considerations in mind. However, priority is handled within the backbone. If congestion occurs, datagrams can be discarded (as in IP) and the end system RTP stack will retransmit as needed. APPN/HPR is designed so that congestion rarely occurs.
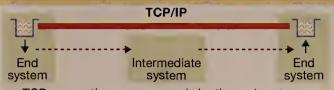
APPN/HPR differs from TCP/IP in that connections are visible only to end systems at two levels: at the RTP layer, where the fundamental connection characteristics are actually implemented, and at the SNA session layers, where SNA-specific connection characteristics are implemented.

APPN/HPR differs from the IP layer of TCP/IP in one important respect. At the time the ANR route is established, a special packet is sent to verify that sufficient resources exist in each node hop-by-hop along the route. If insufficient resources exist, the packet is discarded, and after a time-out, an alternate route will be calculated by the originating node. In this sense, APPN/HPR is actually a hybrid between connection-oriented (RTP) and connectionless-oriented (ANR) protocols; HPR is, in essence, a deterministic connectionless approach. Nevertheless, no session state or flow control feedback is required within the intermediate systems.
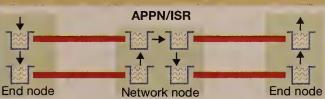
APPN/HPR LINK TYPES

The initial APPN/HPR architecture will define link mappings for frame relay and Logical Link Control 2 local-area networks. In a homogeneous APPN/HPR-only backbone, these two link mappings are sufficient. However, in a heterogeneous backbone with other concurrent protocols such as IP, Novell, Inc.'s Internetwork Packet Exchange (IPX) or AppleTalk, some method is needed in order for the multiple protocols to share the link concurrently. The most straightforward scheme is to define another Point-to-Point Protocol (PPP)-type field and layer APPN/HPR's RTP over PPP.

However, current architecture thinking within IBM leans toward the inverse approach, layering PPP atop ANR (perhaps even atop RTP). This latter approach will be difficult to implement by router vendors. A well-thought-out solution is needed. Convergence of IP and APPN/HPR backbones will be the architectural driver for this requirement, which will be real immediately upon the release of the APPN/HPR architecture.

MAKING THE RIGHT CHOICE

The comparison of connection-oriented and connectionless backbone schemes is a complex analysis since there are many ways of designing and implementing both. Sometimes the performance trade-offs are caused by factors only peripherally related to issues of connection vs. connectionless.

The typical intermediate system today is a router. Routers today have become high-performance packet processing engines with switching throughputs. A significant contributor to router hardware cost is memory — high-speed expensive memory.

Routers achieve their high performance through several schemes:
■ Minimize buffer copying in memory.
■ Implement as much protocol processing as possible for known addresses and paths in fast-path software logic, typically within device drivers directly serving each LAN/WAN port.
■ Commit maximum amounts of processing function to hardware.
■ Use high-speed multiport memory for packet storage.

The choice of a connection-based routing fabric, such as APPN/ISR, has the several implications.

Connections require a higher level of per-packet processing than connectionless traffic, so the router must maintain state, flow control, buffering and pacing on a per-connection basis. The higher level of per-connection processing will be more expensive to implement (and less likely to see hardware implementation) in routers. Routers that forward datagrams usually implement a fast path at or near the device-driver layer (they receive the packet from an inbound queue, perform a quick address lookup and post the packet on an outgoing queue).

Connections also require higher levels of memory (especially because of buffer reservation requirements). The levels of memory required are more severe if buffer pool guarantees are prereserved at connection establishment. (Note that the initial release of APPN routing technology required prereservation of connection buffers for the negotiated adaptive pacing window).

The cost/performance ratio for implementing a connection-oriented router engine is higher than for connectionless backbones. Furthermore, intermediate system APPN/ISR routing, which contains two layers of connection processing (link layer and session layer), will be more expensive to implement and will exhibit lower performance than either IP or APPN/HPR routing.

When applied to high-speed networks, the notion of a connection usually implies two things: the determination of a route and the reservation of resources along the route. All subsequent data flow, however, is datagram-oriented. A given frame or cell can be discarded if network congestion occurs. Our definition of connection differs slightly and includes the function of reliable delivery and flow control on a hop-by-hop basis.

Because of the inherent advantages of connectionless backbone transport for APPN routers, a rapid migration away from APPN/ISR can be expected once APPN/HPR becomes available.

That is not to say that connectionless routing is without its own disadvantages. Because no connection state is maintained, the router has no per-connection method of slowing the rate of datagram arrival during congestion periods. Thus, a connectionless router must ultimately discard datagrams in certain cases. Connection-based routers can apply back pressure (flow control) to the upstream neighboring routers and never be forced to discard a packet.

Also with connectionless routing, problem management and network management are more complex on a per-session hop-by-hop basis.

Newer routing techniques are on the horizon that offer a blend of datagram forwarding services, route setup and reservation services along a given route, essentially providing a hybrid connection/connectionless model.

Migration paths

An April International Data Corp. study attributes positive attitudes concerning APPN adoption to the fact that APPN is the only technology that promises to slowly migrate older 3270 applications to peer-to-peer networking.

APPN prejudice

Corporate sites with a high penetration of LANs are less likely to adopt APPN since they have a sizeable investment in TCP/IP and IPX.

Making the connections
Comparison of TCP/IP and APPN routing techniques.

TCP/IP
End system — Intermediate system — End system
TCP connections are seen only by the end systems. End systems inject datagrams into the network. The intermediate systems forward the datagrams.

APPN/ISR
End node — Network node — End node
APPN/ISR routing is loaded with connection-oriented facilities: Point-to-point links between EN-NN and NN-NN are connection-oriented. For purposes of flow control, each EN-NN and NN-NN hop is connection-oriented.

APPN/HPR
End node — Network node — End node
Like TCP/IP, APPN/HPR provides a connectionless service, although an HPR backbone can prioritize traffic and discard datagrams to abate congestion.

GRAPHIC BY SUSAN SLATER   SOURCE: VEDACOM, SEATTLE

Routt is president of Vedacom Corp., a Seattle consultancy specializing in SNA, TCP/IP, ATM and client/server issues. He can be reached at (206) 527-3434 or via the Internet at tjroutt@halcyon.com.

HOW TO WORK SOME OF THE
TIME
AND STILL LOOK
DEDICATED.
(OF COURSE, WE'RE REFERRING TO YOUR DATA LINES.)

When you need to send data some of the time, but send it fast anytime, use fast packet services via The World Class Network[SM] from GTE. Regardless of your computing platform – mainframe, client/server or a mixed breed – you can send critical bursty data on demand.

With The World Class Network, you can select from a spectrum of fast packet services: Frame Relay, SMDS and emerging ATM. You can also access consultative services and rely on us for all the necessary inter-networking equipment. And through our Network Operations Center, we'll monitor your lines 24 hours a day, seven days a week. Now that's dedication.

No matter when your bursty data starts or stops, fast packet services via The World Class Network can accommodate the workload. In a host environment, Frame Relay transmits data from 56 Kbps up to 1.54 Mbps. MegaConnect™ service, SMDS from GTE, takes the client/server platform to speeds of 56 Kbps to 45 Mbps. ATM transmits voice, data and imaging at speeds of 45 Mbps to 155 Mbps, even if you're using many different systems.

THE WORLD CLASS NETWORK™

Credit card approvals. Hotel reservations. Just-in-time inventory control. No matter what bursty data you need to send, rely on The World Class Network. The advanced communications network that's here today. And only from GTE.

Call for a free copy of The World Class Network interactive game. It will open your mind and your desktop to the power of this advanced telecommunications network.

For more information, call us at **1·800·GTE·4WCN**. We'll fax you back a product sheet like clockwork. Or you can download this information to your computer 24 hours a day by dialing 1-800-483-4927. Call today.

GTE

It's amazing what we can do together.

FAST PACKET SERVICES

IBM exec paints net view of LAN management realm

What does it take to manage local-area networks on an enterprise scope? Network World *recently posed that question and others to Sanjiv Ahuja, IBM's director of enterprise management platforms.*

Ahuja, who is best known for overseeing development of NetView/6000, spoke to Features Editor Charles Bruno last month in Las Vegas at the NetWorld + Interop 94 trade show.

What do customers need today to manage LANs on an enterprise scope?

When we look at management, we have to look at it as primarily two disciplines. One is the very operational management; in a lot of ways it is understood as reactive management. You have a problem and you work on solving the problem a quickly as you know how, with all the tools and services that are available. Then you take it into more of an analytical part, or the proactive part of management, where you look for outages before they happen.

You not only need to know the operational characteristics and any specifics of the environment, you have to know the history of the net, too. A significant amount of analysis and administrative management is involved. Some of the disciplines involved include change management, configuration management, software distribution, asset management. And there's business management accounting, capacity planning, job scheduling, those kinds of things.

On the operational side, there are basic fundamentals — problem isolation, performance and configuration management. Then there are disciplines which span both of these areas; they are things like security and automation. Automation is used most today in traditional host operating environments. But over time, the LAN environment will require the same type of productivity and scheduling flexibility.

I would also venture to say that customer needs and requirements from a business perspective are no different than they were in a centralized environment. Customers expect 24-hour, seven-day-a-week operation of their LAN infrastructure. Whether it's distributed, client/server, centralized — they don't care. They want access and the ability to use that infrastructure in the data that is in their environment at any minute, any instant of time, from anywhere globally.

Those requirements are not really different. However, the complexity of the LAN environment is a lot more than a centralized one. Because now you are looking at an environment where the data is distributed, your applications are distributed, so you're electronic infrastructure instead of being a very controlled centralized infrastructure, has really spread out and become a very horizontal infrastructure. To manage a horizontal infrastructure is a lot harder than managing a centralized infrastructure. There's no single protocol. The level of complexity is greater; the impact on personnel resources is greater than a centralized environment because expertise is spread around.

> "Customers expect 24-hour, seven-day-a-week operation of their LAN infrastructure. Whether it's distributed, client/server, centralized — they don't care."

A lot of what we have learned in centralized management is very, very valuable to LAN managers. The basic customer requirements for availability, control, automation are very much the same and increasing.

Are the tools — not just IBM tools but LAN management tools in general — mature enough today to provide users with the level of distributed management necessary for enterprise management?

Many of the tools are there, and many are not there. There are disciplines in which steps need to be taken.

Operational management is something we've been doing for years at IBM, largely from a centralized NetView. I will give you examples of why that vision maps into a distributed environment. Take NetView/6000, for example. You can have geographically dispersed local monitoring systems, like systems monitors and other products, in different geographic locations sending exception information to a centralized NetView. There are multiple NetViews spread in the enterprise that are feeding information into a host NetView. That all exists today. That's for operational purposes.

However, we are taking steps on the next level to provide cross-platform functionality in problem management, change/configuration management, access control, software distribution and security.
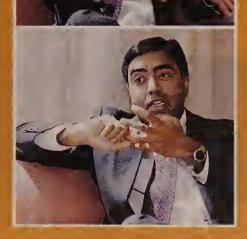
When you really look at it, each management discipline has to be addressed, because when our customers roll it out in their enterprise, they want full flexibility and choice. They want to do operational management in a largely distributed way, centralize the rest of the management, change configuration, asset management in one location — and integrate security across all of those areas. And tie it in with what I have running on the System 390 today. That's right, and it's very fair. We have to step up to it. That's why our investment in building synergy between NetView/6000 and System 390 is the correct approach. Just look at what kind of information gets exchanged — problems, performance, configuration — between the platforms in either direction. We are looking for an even stronger level of interoperability.

A recent survey *NW* conducted of Simple Network Management Protocol-based management systems finds that only about two-thirds of respondents were managing 500 elements or less. Are customers in Unix environments using SNMP-based managers to control LANs across their enterprises, or are they just managing small LAN environments?

The majority of customers using Unix-based management systems are using them to manage TCP/IP environments. There are some customers that are beginning to look at integration of other environments, and then there are customers that are integrating other environments.

There are a set of reasons for that. One is that customers haven't had the tools to integrate these different environments. They may have a TCP/IP environment and an IPX environ-

Corporate users can borrow heavily on the lessons learned in host-based net management, says IBM's director of enterprise management platforms.

ment, for instance. This is very common in a distributed LAN environment — and maybe it supports SNA, too. Suppose a workstation with an SNA address and a TCP/IP address fails. There is no management system today currently available that can make the correlation between these two addresses. We shipped framework services in the middle of last year with NetView/6000 Version 2. But the level of integration we are providing now is through SNA Manager 6000. That rolls out this month.

Now, you would ask, 'Is it realistic to move the management of a 50,000 logical unit network onto a Unix system?' Absolutely not. Not in the foreseeable future.

If you really want to do some very interactive, active management of your SNA network, a smaller SNA network is realistic. But you may want a window into a subset of a large network integrated with that local-area network. That level of integration is [provided by existing technology] today.

Suppose you want to get to the next level, you want to do complete enterprisewide management. Then you are going to take another view of it. But to do real interactive management of 100,000 [logical unit] networks from a Unix system is unrealistic.

So that's still the role of a host-based NetView and always will be for the foreseeable future?

Yes, for the foreseeable future. To expect a single Unix machine to manage some of our large customers' nets with billions of logical units is absolutely unrealistic.

Isn't NetView/6000 capable of managing a very large enterprise network from several workstations?

We just took a huge step here. You need some core services which are completely distributed. You can't have a centralized view because one of the requirements a customer has is, 'Give me a centralized view of my complete SNA network, but instead of providing it on a 390, view it on a single Unix machine.'

There is no Unix machine large enough which would let it happen. However, there are some other issues. A lot of host NetView usage is for automation. One of the fundamental rules of automation is to do the automation closest to the source of the problem. It would be economically unwise if you start moving all the automated procedures for collecting data across a channel or a transmission link to a System 390 and do the automation there. If you started doing automation

> "Now, you would ask, 'Is it realistic to move the management of a 50,000 logical unit network onto a Unix system?' Absolutely not. Not in the foreseeable future."

of all Unix machines centrally on a 390, you'd spend too much on bandwidth and too much on your CPU usage. So a lot of that automation will probably never move. It's technically the right way to do it. Whatever your resources, if it's AS/400, Unix, Windows, do it at the source of the problem. Don't move the information around if you don't have to, because moving the information costs you in bandwidth, and there's a time lag and you don't understand that environment as well if it's not local, right there.

So that says you are retaining a lot of post-NetView automation services. Now you could envision a scenario that's not unrealistic at all where you have a completely distributed NetView/6000 where your topology is millions of [logical units] in a completely distributed database, spread all over the enterprise. We are not that far away from that kind of scenario, but manageability of that database itself becomes a very tough task. It would be much easier to manage it on a single system. When you spread it, you have to keep it replicated. So there are some technology barriers, but they are not unsurmountable, but then they lead to a different set of issues.

We spoke about host-based NetView and NetView/6000. Are there customers out there that have resisted the move to NetView/6000 and would want a smaller product, say NetView on OS/2?

Absolutely. There are customers that have skills in OS/2; they made a significant corporate-level commitment and frankly see no business value in adding a Unix system to their environment. They want to do management of their LANs, systems management of Unix systems, DOS, Windows systems, NetWare servers, as well as a level of network management for hubs, routers, etc. And we see that as a very strong requirement, so sometime this year, we are going to roll out the next iteration of LAN NetView technology or NetView for OS/2,

Freedom of choice was never this fiscally rewarding.

Introducing TCG. The *other* local phone company.℠

As far as most people are concerned, for local telecommunications services, there's been the Phone Company and there's been the Phone Company. But now there's also Teleport Communications Group (TCG).

Saving money with TCG.

By choosing TCG® services, your business will be rewarded with immediate and substantial savings on local service. Whether it's to access your long distance carrier or your office across town, TCG provides a full range of dedicated circuits.

Reliability has its own rewards with TCG.

TCG's self-healing fiber optic network and diverse routing allow the TCG network to offer 99.99% uptime: reliability pure and simple.

Not lip service. Real service.

Real customer service. TCG installations and repairs are executed within hours, not days. And TCG's customized applications are a true reflection of your business needs.

Call TCG. The nation's pioneer in alternative local telecommunications.

1-800-969-5686 lets you find out even more about TCG. For over a decade TCG has been the leading alternative in enhanced voice, data and video communications. And TCG currently serves over 150 communities nationwide on its ever-expanding fiber optic network.

So for maximum savings and enhanced local telecommunications service, call TCG. It's a very rewarding choice.

CALL 1 800 969 5686

TCG

which will essentially provide that solution to our customers. Then you can have a very synergistic management scheme with NetView running on whichever platform is most appropriate for an individual company.

Users complain about the lack of integration between network management platforms and the applications that run on them. What is IBM doing to improve integration between its management platforms and the applications that it brings to customers?

Let's look at the various levels of integration. There's a visual integration that is a very first step in integrating applications with a management framework. The second step is integration of the information which relates to topology and other things. The third step is integration of the management data itself. So you can use the data for operational reasons as well as analytical analysis. You look at those two or three things, and there are integration points.

It's the management platform vendor's responsibility to provide tools and services to do visual integration. It is also management platform vendors' responsibility to provide tools and services which do topology integration.

It is also management framework vendors' responsibility to provide tools which perform data integration.

So technology has to exist in the framework that lets this integration happen. On Unix systems, besides visual integration, that sort of integration just hasn't existed. We were the first vendor to provide tools for topology integration in the middle of last year. Most of the Unix systems today don't even provide an open database from which to access any data. If they do provide a database, it is limited to topology information, if that. And that is the best anybody provides. Even if it is that, it's not open. It's not accessible.

Some of the things we are doing is making NetView/6000 available on the largest set of hardware/software platforms in the industry. We are on AIX today. We have shipped it on Sun. We are working with Digital. We have shipped it on Alpha OSF/1. We will ship it on OpenVMS. We will ship it on NT, on Intel and NT on Alpha.

I just mentioned we are taking those [application program interfaces] over to OS/2. So if you think about it, it has the broadest set of hardware/software platform choices today, with one set of APIs.

But that's not the same as the industry adopting a common API set.

I think vendors would like a consistent set of APIs. But a consistent set of APIs should not imply a consistent set of functions. As you add function, you have to add APIs. And when you add APIs, if the rest of the folks in the industry do not catch up with that vendor, they start lagging behind.

I will give you a specific example. Nobody provides topology integration today besides IBM.

We are beginning to go into beta on NetWare 6000 Version 3 that will give users a choice of databases and complete access to the information. Now the rest of the industry hasn't caught up yet. So we are actually very willing to license this kind of technology — of course, at reasonable terms.

The nice thing would be for us all to stop and wait for the others to catch up. Well, our customers' needs are not stopping. We have to continue to roll out new functionality. However, it is very critical that we keep it open. We use industry standards where they exist. We provide open APIs. And we provide help and assistance to the vendors to use those APIs.

IBM offers a large number of net management products. Customers seem confused about which products to use. Is there any consolidation effort under way to simplify IBM's line of management offerings?

Let's look at two sets of customer needs. One is in the framework platform, and the other is in the application arena. You should look at IBM's framework direction being NetView.

In the application arena, there are a lot of choices. We have a job scheduling application, we have a load balancing application, we have an accounting application, we have trouble-ticket applications on the host system.

Does that cause confusion? Two things. One is that customers have asked for flexibility in choice. When you have a very flexible solution with a lot of choices, it could make you question if you are making the right choices. The right choice is what suits that customer's needs, not what a vendor's product limitations are.

There are customers who do have multiple choices. In those environments, we just need to work very closely with them so they understand the choices that best suit their needs. ◼

"Send Three and Fourpence, We're Going to a Dance"

BY SCOTT McCREADY
International Data Corporation/Avante

As the old British joke goes, a daring group of soldiers during World War I sent an urgent message to Group HQ: "Send reinforcements, we're going to advance." Technology being limited to messengers that risked their lives sprinting from one trench to the next, the message eventually reached Group HQ, where it was finally interpreted as: "Send three and fourpence, we're going to a dance."

As advocates for reengineering continue to advance their theories of mass destruction, you hear a great deal about the glory of reengineering but very little about the risks or the rate of recidivism associated with it. Reengineering has become an almost patriotic call to arms, but victory is never guaranteed. Nevertheless, the reengineering movement has fulfilled one critical role: it has become a battle cry exhorting: "all is not well in US business." However, while the re engineering banner is still flying high, it is an inappropriate rallying point for most companies.

Instead, IDC/Avante believes that the concept Work Management will be the agent of change for US corporations. Work Management does not require a leveling of an organization's existing structure, business systems and technology investments. Most organizations cannot afford this degree of change. Instead, Work Management recognizes the investments of the past, the current level of expertise in the organization and then looks to the judicious application of new technologies to provide quantum leaps in productivity.

Work Management is the coming together of many technologies including: work flow software, imaging, document management, groupware, networking, e-mail, folder software, and products such as Lotus Notes in such a way that they make a definitive contribution to the profitability of the organization. In fact, IDC/Avante recently helped more than 65 companies in 12 different countries quantify their investments in Lotus Notes and found an average three year rate of return of 179%.

> **Work Management recognizes the investments of the past, the current level of expertise in the organization, and then looks to the judicious application of new technologies to provide quantum leaps in productivity.**

Increasingly, IS will be required not only to achieve these types of results but also to have the skill-set to document the financial impact of all types of information systems.

Recognizing that Work Management will require an amalgam of new technologies to automate business processes on an enterprise-wide scale, IDC/Avante has designed a three-day conference to provide the necessary education for IS and Networking professionals to become proficient in the field of Work Management. Join us to learn how you can implement information technologies that will lead to substantial rationalization of your business systems and how you can measure the financial impact.

If you are an IS manager or a networking professional whose voice may be heard but not always correctly interpreted in the trenches, bring your favorite project champion and attend Work Management '94: The Conference on Business Process Automation.

Scott McCready is a principal at IDC/Avante Technologies, a Framingham, MA market research firm that specializes in providing expertise on workflow, imaging, and related technologies to vendors and IT customers. The company recently completed a landmark study *Lotus Notes - Agent of Change: The Financial Impact of Notes on Business.*

Dr. Sigmund Freud

Dr. Sigmund Freud

Announcing dual analysis for ATM.

The new ATM capabilities on our DA-30 Internetwork Analyzer are indispensable to anyone developing ATM internetworking solutions. With dual simultaneous analysis in a powerful, easy-to-use Microsoft® Windows™ environment, the DA-30 provides the tools you need to perform stress testing, traffic analysis and benchmarking across your entire internetwork. Today, DA-30 broadband capabilities support HSSI and DS3 ATM. Tomorrow, new ATM interfaces will continue the DA-30 tradition of providing flexible internetworking solutions. But then, what else would you expect from Wandel & Goltermann, the developers of the new IDMS remote monitoring family and many other internetworking analysis tools?
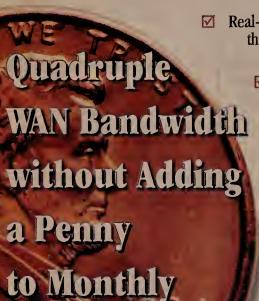
For more information and the phone number of your local sales office, call us direct. Because when it comes to internetworking analysis, no one can provide you with better solutions than W&G.

Wandel & Goltermann
Electronic Measurement Technology

Dark past

gral role in carrying the museum's message, is considered by staffers to be as mission-critical as any business net.

"You only have one opportunity to show your information to people," said Yechiam Halevy, a systems integration consultant who helped build the multimedia-based local net. "They won't come back a second time if the system goes down while they're visiting."

The showcase of the museum's networking investment is its Wexner Learning Center, which is open to visitors and the museum staff. This is a separate area off the main exhibit where 30 touch-screen IBM ValuePoint 486 workstations provide on-line access to a NetFRAME Systems, Inc. NF450 superserver running Novell, Inc. NetWare 3.11 and housing more than 16 hours of full-motion video, 6,000 photographs, six hours of music, 700 maps, plus the text of the MacMillan Encyclopedia of the Holocaust.

When a visitor sits down at one of the workstations, a graphical front-end application presents a variety of topics to choose from. Visitors can access various types of information simultaneously, such as both a map of post-World War II Germany and a video clip of a Holocaust survivor interview.

"This is a totally on-line system, all on the hard drive of the server," said Arnold Kramer, director, technical services for the museum. "Now the Net-FRAME has about 60G bytes of disk storage, but we expect it to double in size rather quickly."

NO SIMPLE TASK

The museum developed its multimedia applications in house, which Kramer said was no simple task.

"We had two teams, one that developed the software and put the network together, and another that developed the content," he said. "The technical team had about six people on it, and the content team had 12 or 13 people. The whole project took about two years."

Beyond the Learning Center, Kramer said the museum is planning this summer to publicize the address of its existing WWW server to give Internet users direct access to museum resources. The museum uses a dedicated Unix-based machine running Mosaic as an Internet server.

Once the address is made public, users will have access to general information about the museum. But the next step is to give users access to the museum's databases of historical and other Holocaust-related information.

BRINGING IT TOGETHER

The Wexner Learning Center and Internet servers are both part of the museum's day-to-day network. This is a 300-node network with about seven servers separated into about 11 token-ring segments running SynOptics Communications, Inc. hubs.

The servers run NetWare 3.11 and support Transmission Control Protocol/Internet Protocol. The museum staff manages the net using SynOptic's Optivity software via a Sun Microsystems, Inc. SPARC-based machine.

The seven servers and the superserver are connected via Fiber Distributed Data Interface backbone. Until just last month, the museum was running traffic at 16M bit/sec across a fiber-optic token-ring backbone.

"The reason for the FDDI backbone was to let us ship video around the network," Kramer said.

Most of the seven servers run administrative applications — such as databases, word processing and electronic mail — for the museum staff. But more staff members are requiring video to the desktop, and many of them need access to the NetFRAME server's constantly changing information.

The combination of this and "the pressure to propagate Windows throughout the net," Kramer said, forced the move to FDDI for the additional bandwidth.

"Before, the performance just wasn't good enough," he added. ⊠

Mail mgmt.

Link Services and author of *Managing Messaging Networks*. "We want to keep it open enough that we can keep adding different applications or modules to this."

The EasyLink framework lends itself to deployment in large, distributed corporate environments with equipment, applications and messaging systems from multiple vendors, Ananthanpillai said. And AT&T plans to roll out a management service, based on the framework, that will allow EasyLink customers to perform administrative functions from a customer premise-based graphical interface, he said.

Ananthanpillai said he didn't know when the service would be rolled out.

The EasyLink framework is based on AT&T's Unified Network Management Architecture (UNMA), a 7-year-old blueprint for centralized management of global voice and data networks.

But the EasyLink framework strays from UNMA's dogma of relying on the Common Management Information Protocol.

Indeed, the centralized manager of managers in the EasyLink network is one of the leading SNMP platforms — Hewlett-Packard Co.'s OpenView. A user at a single OpenView console at the EasyLink operations center in Bridgeton, Mo., can delegate management tasks to at least 10 element management systems distributed throughout the global network, organize alarms coming in from those systems and respond to alarms on behalf of the element managers.

The element managers include products for managing Tandem Computer Corp. and Unix systems, which serve as EasyLink MTAs; Unix workstations and personal computers, which serve as mail gateways; front-end processors, Wellfleet Communications, Inc. routers, AT&T Datakit switches, Optical Data Systems hubs, and RS-232 and ASCII devices.

All this gear makes up an EasyLink node complex, of which there are more than 35 across the globe. Node complexes handle the switching and service provisioning for the EasyLink network.

Though the view of the network is centralized, management is distributed. Through automated knowledge-based routines and applications, the element managers play a big role in correlating and responding to alarms and alerts from all the devices and software that handle messages going through the network.

But much of the management is performed locally on EasyLink nodes.

AT&T has developed and patented so-called local intelligent monitors that perform systems and application performance monitoring, alarm filtering, configuration management and event thresholding.

Using the AT&T developed MIB, which defines all the management attributes and variables of EasyLink's MTAs and messaging gateways, the intelligent monitors can track activity in message queues in the MTAs.

For example, a delivery queue filling up could prompt the monitor software to issue an alert to the local element manager. Using its knowledge base, the element manager can then check the messages to determine if one is corrupted.

The element manager can then remove the corrupted message, place it in temporary memory, correct it and add it back into the queue.

If problem determination or resolution extends beyond the element manager's knowledge base, the element manager will alert the central Open-View console to intervene.

"We've developed hooks in the network that detect potential problems early enough that it gives you a warning to react to before the customer knows about the problem," Ananthanpillai said.

The EasyLink net operators can prioritize alarms, control the rate at which they are issued and the amount of text in the alarm message, he said.

The element managers use expert system-based fault, configuration and performance management applications that will be able to infer from alarm messages what the most likely cause of a problem is and what the solutions are. The inference engine is based on technology from Gensym Corp.

The knowledge-based applications share a common Sybase, Inc. data repository so they can access the same routing, asset, configuration, physical topology and connectivity information. ISICAD, Inc.'s Command system provides EasyLink's asset, configuration and physical topology management capability.

"The goal is to consolidate as much information as possible and correlate it so that the operator is given the right information," he said. "They don't have to chase ghosts." ⊠

For more info...

AT&T's Ananthanpillai can be reached at raja@attmail.com. *Managing Messaging Networks* is published by Artech House, Inc. in Norwood, Mass. It costs $62, or $53 if ordered before Aug. 31. Artech can be reached at artech@world.std.com.

Raid

The rerouting happens to a degree today in "bypass" situations, where users circumvent the LEC with a private line for long distance. "But it happens less than you'd think," said Eric Paulak, associate publisher for the Center for Communications Management Information, a provider of rate and tariff information in Rockville, Md.

A business user could also route local toll calls to the AT&T network by dialing AT&T's access code. However, the switch modifications would make that access transparent to end users.

As of July 1, **45 of the 48 contiguous states will allow competition for local toll traffic.**

The holdouts are California, Nevada and Virginia. In Arizona, Arkansas, North Dakota and Utah, carriers must seek approval to compete against the local exchange company.

But this does not mean all AT&T customers will jump aboard. Ken Fauerbach, manager of user services at New York University, an AT&T PBX shop, said: "AT&T is going to have to do significantly better on its local toll rates to make [reprogramming] worth our while."

Intra-LATA toll competition is a big, new threat for the local telephone companies that began last September. Since the 1984 AT&T divestiture plan, intra-LATA traffic had been handled almost exclusively by local carriers, while long-distance carriers only carried calls crossing LATA boundaries.

While local carriers are facing growing challenges to their local loop strangleholds, they are still restricted from participating in the long-distance market. Legislation being considered by congress could change that if passed this year, however.

COMPETING WITH RBHCS

The PBX efforts go hand in hand with consumer-oriented programs from AT&T for capturing local toll business. The carrier announced last month that it would charge consumers 10% less than regional Bell holding company Ameritech for calls of 15 miles or more within the greater Chicago area. Callers have to dial the five-number sequence that automatically routes them to the AT&T network.

AT&T is launching a similar program today in Missouri.

To date, any Definity Generic 3 user could have opted to program the switch with the 1-0-XXX digits that route calls to secondary carriers using AT&T's Auto Route Selection software.

But the new program would alleviate PBX managers from undertaking complicated programming in states that have gotten rid of the "1+" differentiator between local and toll calls.

States that have done away with the 1+ scheme include Illinois, New Jersey, New York, Pennsylvania, Virginia and West Virginia.

The AT&T efforts are among several afoot by the carrier to extend its reach into the "last mile" of user networks. Another is a new fiber ring access service AT&T is delivering via partners that include the LECs (see story, page 1) and AT&T's pending merger with wireless giant McCaw Cellular Communications, Inc., slated to close in August.

AT&T wants more control over end-to-end network reliability and costs for its customers' networks. But it is also being lured by the lucrative nature of the multibillion-dollar intra-LATA toll business and driven by cutthroat competition from the likes of rival competitor MCI Communications Corp., which is building its own local infrastructure under its MCI Metro venture. **Z**

Comments?

Touchdown

formation Exchange Microsoft user conference here. They may be disappointed, however, as many also expected the announcement to come at the Electronic Messaging Association's annual conference in Anaheim, Calif., in April.

Bill Sornsin, project manager for Microsoft's business systems division, said users at the conference will see that Microsoft is "putting together the best of groupware and the best of messaging." Microsoft has designed a client/server E-mail system with an object-enabled database for information sharing.

Digital Equipment Corp.'s Linkworks and Lotus' Notes will be the top groupware competitors to Microsoft's offerings, while Microsoft plans to target Digital and Hewlett-Packard Co. in the E-mail market.

Larry Moore, vice president for telecommunications services at Lotus, said Microsoft Exchange remains more of a promise than a competitive product, he said.

"The word 'smoke' comes up over and over again" in discussions with Notes customers who have looked at Microsoft Exchange, Moore said. Novell, Inc. is actually the company that should really worry about Microsoft's messaging products, he said, since Microsoft will likely try to use Touchdown as a Trojan horse to get its NT operating system into NetWare shops.

Meanwhile, users and analysts are taking Microsoft Exchange seriously.

John Lisiak, senior software engineer at Chevron Information Technology Co. in San Ramon, Calif., said his company will test a beta version of Microsft Exchange "to decide if it's a way to pull back from Notes."

Lisiak expects to get a test version of the product from Microsoft later this month; Microsoft launched its first wave of beta testing last November.

Lisiak said Chevron has only 1,500 Notes users, but he prefers a good messaging system with some Notes-like information-sharing capabilities, such as those Microsft Exchange promises. Notes focuses on collaborative groupware functions, downplaying the role of messaging, he said.

Microsoft Exchange's support for Transmission Control Protocol/Internet Protocol also means that Chevron would be able to link its 20,000 client/server and host-based E-mail users directly through TCP/IP, Lisiak said. This could enable Chevron to eliminate the 57 Professional Office System (PROFS) gateways it maintains worldwide, he said.

Michael Goulde, senior analyst at Seybold, said Microsoft Exchange compares favorably to Notes in some areas. Its support for TCP/IP's Simple Mail Transfer Protocol and other mail transfer protocols means the product can handle more document types than Notes, he said. And future versions will give it comparable application-development features.

"But it's not clear that [Microsoft has] the quality of replication that Notes has," he said.

Sornsin says Microsoft expects to ship the product by year end. **Z**

Ring

the regional Bell holding companies, GTE Telephone Operations and cable television company Time-Warner, Inc.

AT&T said it has tested rings in Colorado, Georgia, Nebraska and Texas, states where the service is now available. The company will work with customers and other carriers to roll out the capability in areas where fiber rings are not already in place.

THE LAST MILE

One of AT&T's major goals is to extend its service reliability to the last mile of a wide-area network, said Hemant Vaidya, AT&T product manager. In addition, AT&T is looking to pare customer costs and become a central source for network sales and service.

AT&T will have monitoring access into the rings so it can manage the network for customers end to end. This includes extending customer-contracted network uptime guarantees across the rings.

Pricing will be determined on a customer-by-customer basis.

"In theory, if this works, customers wouldn't have to buy diverse local loops," said Christine Heckart, senior consultant at TeleChoice, Inc. in Verona, N.J. "That offers users the opportunity to eliminate headaches

> **"In theory, if this works, customers wouldn't have to buy diverse local loops,"** Heckart said. **"That offers users the opportunity to eliminate headaches and save money."**

and save money."

She said the cost of using geographically diverse local loops is at least double that of purchasing an end-to-end network from a single service provider.

Rick Malone, principal at the Vertical Systems Group consultancy in Dedham, Mass., said, "In a way, you could look at this as overdue for large users from AT&T. This is the way everybody wants to go."

AT&T customer Norwest Technical Services in Des Moines, Iowa, for example, sought out its own self-healing network about a year ago from the CAP Midwest Resources, Inc.

"At the time, AT&T didn't have any of these agreements, so we went outside to find our own provider," said Eric Shafer, communications analyst.

The plan has saved Norwest about $2,000 per month, and network uptime has jumped from about 85% to 99.5%.

But Shafer said he would prefer to have one vendor "to get the whole thing taken care of."

MCI Communications Corp. has been working with CAPs and local exchange carriers for about a year to provide MCI Custom Access, which is similar to Accu-Ring, and can support half a million buildings nationwide.

Neither Sprint Corp. nor WilTel have formal programs, but they will build such a network for customers on an individual basis. **Z**

Notes

tween Notrix and other Notes development tools is its use of an event manager that lets developers easily build event triggers into their applications. This would let an event within a Notes database, such as an end user adding data, trigger an external application.

The company last week also announced a related product called Notrix Composer for EDA/SQL. It allows Notes users to access back-end SQL databases via Information Builders, Inc.'s EDA/SQL Server, a database gateway system.

Robert Wesneski, Notes administrator at Phillips Petroleum Co. in Battlesville, Okla., praised Notrix's ease of use.

He said its reliance on REXX allowed him to easily build a beta application for linking problem-tracking databases both in Notes and on an IBM MVS mainframe.

Users will also have the option of going directly to Lotus for Notes application development software. Lotus'

> **Notrix notes**
>
> **Percussion Software's Notrix is available now starting at $3,495.**
>
> —:—
>
> **Notrix Composer, which includes Notrix, costs $4,995.**
>
> —:—
>
> **A developers' kit is available for $995.**

ViP will be aimed at development of client applications and interfaces, which are mainly what Notes developers are calling for, said Alex Neihaus, senior product marketing manager for Notes ViP at Lotus.

ViP, a graphical user interface-based development tool, will be used to trigger external applications and give Notes users access to SQL databases, Neihaus said. It also incorporates tools for extracting Notes data for reporting and analysis, he said.

The initial version will be limited to Windows platforms, but Neihaus said Lotus is looking to extend it to the other platforms on which Notes runs. Pricing was not announced.

Lotus this week will also announce a new version of its Approach desktop database that can be used as a front end to Notes for analysis and reporting.

Version 3.0 will include access to Notes via Lotus' FX technology, which lets Notes exchange data fields with desktop applications. It will ship in July and cost $129.

©Percussion Software: (617) 267-6700; Lotus: (617) 577-8500.

AT&T to help users sidestep net outages

BY JOANIE WEXLER

Basking Ridge, N.J.

AT&T is expected to announce today an enhancement to its private-line management system that predicts network disruptions, allowing net managers to head off downtime at the pass.

The new Windows-based Proactive-Predictive-Preventive (P-P-P) application builds on AT&T's Accunet Information Manager (AIM) tool for real-time monitoring of T-1 as well as fractional T-1 and T-3 links, according to analysts briefed on the announcement.

P-P-P moves beyond AIM's ability to merely alert users of a disruption. The new application will continually analyze conditions, parameters and historical data to predict deterioration of a link, thus allowing AT&T or the user to prevent more outages, analysts said.

In addition, P-P-P will allow users to issue a trouble ticket electronically; AIM users must do this manually today. This feature should expedite the process of line restoral, users said.

"[P-P-P] is absolutely needed," said Charles Arnett, manager of network services at Accunet shop LSI Logic, Inc. in Milpitas, Calif.

Even though LSI has an elaborate command center, Arnett said the company sometimes finds it difficult to determine the cause of network problems. That's because LSI does not have visibility into its carrier's network from a quality-of-service standpoint, he said.

> **P-P-P will allow users to issue a trouble ticket electronically; AIM users must do this manually today.**

For existing AIM users, AT&T will reportedly provide the predictive service for free. For new AIM users, AT&T will assess a service charge of 2% to 3% of the cost of overall Accunet service, analysts said.

An AT&T spokeswoman acknowledged the upcoming Accunet announcement but declined to discuss details.

Steven Laak, senior network planner at General Mills, Inc. in Minneapolis, said anything that improves restoral time is worth exploring.

"The one thing we don't want is to run redundant links everywhere, because what we're doing [in that case] is safeguarding ourselves against [AT&T]," Laak said.

The predictive scheme could minimize outages and the need for elaborate redundancy, he said.

AIM has ties to AT&T's umbrella Accumaster enterprise manager and its Fast Automatic Restoration restoral system, which guarantees a maximum net restoral time specified by the user.

In its initial version, P-P-P is still proprietary to AT&T under Accumaster, but it will eventually move to Simple Network Management Protocol platforms for blending in with customer data communications networks. One analyst estimated that the port would take one to two years.

→ Senior Washington Correspondent David Rohde contributed to this report.

Report

ment & Budget (OMB) to publish an annual document offering "a single, integrated view of agency internetworking and interoperability across the government and the public."

"I asked that this panel be set up," said Bruce McConnell, OMB chief of information policy. "We'll do whatever the panel recommends."

Jim Flyzik, telecommunications director at the Department of the Treasury and chairman of GITS, said the group — which claims representatives from about a dozen federal agencies — is willing to carry out the report's recommendations.

Roger Cooper, information resource management director at the Department of Justice and a GITS member, added that the group is committed to the administration's "reinventing government" goals. He added that GITS intended to issue a "vision paper" by the end of the month.

However, the decision to turn such a huge coordination task over to the dozen network managers in the GITS group, who all have full-time jobs at their own agencies, could be risky in the long run.
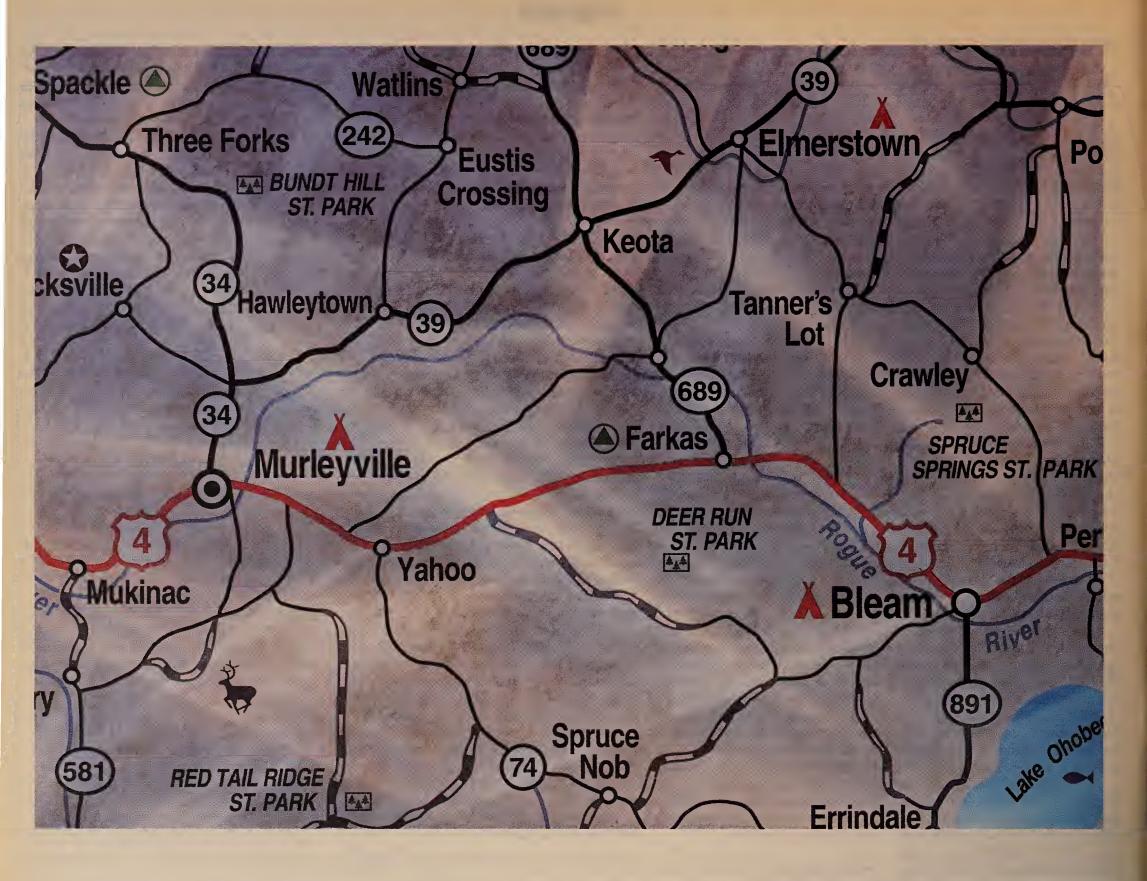
Mel Cooper, president of the systems group at Computer Sciences Corp., a large systems integrator with years of experience in government project, noted that the nature of government is that people frequently change jobs.

In a speech he gave at last week's Armed Forces Communication and Electronis Association show here, he pointed out that the administration must eventually find a way for its efforts continue on through the inevitable string of successors.

For example, the IITF and its GITS working group were created as ad hoc organizations at the wish of the White House. If the White House changes hands in the next election, the future of such efforts could quickly evaporate.

"We do need to worry about this in the future, but not right now," McConnell acknowledged. "We have at least two more years before that."

Attack

together the enterprise by better positioning the technologies, such as ATM, APPN, HPR and AnyNet, at its disposal," said Frank Dzubeck, president of the Communications Network Architects, Inc. consultancy in Washington, D.C.

NPM Version 2 Release 2 will let users monitor NetWare LANs and Systems Network Architecture/LU 6.2 sessions. NPM today collects SNA/3270 and Token-Ring LAN performance data from host VTAM and the Network Control Program (NCP) on the front-end processor.

In order to grab data from the NetWare server, users will need NetWare for SAA or the NetWare Management Agent for NetView, available from Novell.

With the new NPM, users will be able to spot bottlenecks and monitor CPU utilization, server volumes or load balance for heavily used NetWare or LU 6.2 nets.

"Our next focus is to bring NPM to NetView/6000 so users can monitor SNA, Token-Ring and Novell environments from there," said Harry Hall, IBM's NPM product advocate.

NPM 2.2 will be available June 24 for $52,800, plus $14,400 for the Novell management capability.

IBM will also announce NetView Distribution Manager (DM)/6000 Version 1.1, which increases the number of clients controlled by a single server and supports software transfers among multiple servers. NetView DM/6000 resides on a RISC System/6000 server and lets users distribute software and data to attached clients.

NetView DM/6000 1.1 increases the number of clients one server can update from 117 to 400. It also lets users designate a single box as a central distribution server for multiple NetView DM/6000 servers.

Users have the option of going server-to-server over Transmission Control Protocol/Internet Protocol or SNA/LU 6.2. Server-to-client communications must still be done over TCP/IP. NetView DM/6000 will be available June 30, starting at $900.

Also opening up a variety of new configuration options, IBM will add support for APPN Network Node (NN), frame relay lines and Dependent LU Requestor (dLUR) to its 8250/8260 LAN hub lines. APPN NNs control routing and directory functions in APPN nets, while dLUR lets 3270 data flow over APPN links.

APPN and frame relay support more tightly integrates the hubs into IBM's enterprise backbone strategy while boosting interoperability.

NEW FOCUS

In what is being called a refocusing of strategy rather than a product rollout, IBM is expected to define a new multiprotocol role for its HPR technology. Positioned as the "real APPN," HPR promises to eliminate problems with APPN by adding dynamic rerouting around failures, better performance, improved data flow control and class-of-service features.

But IBM wants to take HPR beyond APPN and enable it to handle TCP/IP and other protocols.

While IBM would not confirm it, sources indicate that the company may pull this off by tying HPR in with its AnyNet technology. AnyNet is an implementation of IBM's Multiprotocol Transport Networking code, which makes applications independent of their underlying network. IBM already supports AnyNet in its OS/2 and MVS systems. AnyNet for VTAM is expected within the next couple months.

"We have SNA and are quickly moving toward a multiprotocol network, and the two just don't marry very well," said Tom Donell, director of network services for Phillips Petroleum Co. in Bartlesville, Okla. "If AnyNet allows us to have 3270 sessions coexisting with TCP/IP, we'd [be interested]."

"HPR and AnyNet will provide IBM the fast and reliable channels SNA and APPN sessions need with the multiprotocol support required to compete in the multivendor world," said Anura Guruge, an independent analyst based in New Ipswich, N.H. "Combining these technologies is IBM's master plan for [supporting] legacy applications over multiprotocol backbones."

David Passmore, president of the Decisis, Inc. consultancy in Herndon, Va., agreed but said IBM faces a long, hard fight to implement it.

"You're talking about two technologies that have a ton of promise but aren't in the marketplace yet," he said. "IBM has to get AnyNet and HPR on its own platforms and convince other vendors to do the same."

On that front, IBM will soon deliver AnyNet support for Windows and AIX environments, analysts said.

IBM is expected to refresh its ATM strategy, which will include some new products such as its 25M bit/sec ATM adapter card, and a more clearly defined role for its ATM switch line (NW, May 9, page 1).

Details were sketchy, but analysts said IBM's Transport Network Node ATM switch will come in at about $2,000 per port, which will make it competitive with switches from vendors such as Fore Systems, Inc.

Product rollout call

- AnyNet ports to AIX and Windows platforms
- AnyNet support for NETBIOS Extended User Interface
- ATM adapters
- HPR extensions for AnyNet
- HPR for 3172 and 6611
- IBM 8250/8260 hub software with APPN support
- NetView Distribution Manager/6000 update
- NPM upgrade
- Transport Network Node ATM switches

Same-day service. Anywhere in the country.
For three years. At no extra cost.

PC Server
Model 95

Mission-critical PC Server 85/95 service from IBM.* Getting it
from anyone else may be mission impossible.

IBM®

Cast your vote today in our Enterprise Technology Awards!

Dear *Network World* Reader:

Today, you're building the enterprise networks that support a new generation of client/server and distributed applications — applications that can give your organization a competitive edge in a tight marketplace.

Network World's third annual Enterprise Technology Awards program offers you the chance to recognize your strategic partners — the vendors delivering the interoperable products and services that support your enterprise network.

Attached to this week's issue is an easy-to-complete survey form listing a variety of product and service categories. In each category, simply select the one vendor whose products you think are the most interoperable — those that work best in a multivendor enterprise network environment.

If you are missing an entry form, please call us at (800) 622-1108 for more information. All entries must be submitted by July 1, 1994.

The winners of the 1994 Enterprise Technology Awards will be featured in a special issue of *Network World,* and the awards will be presented at NetWorld + Interop in September.

Unlike other publications, *Network World* relies on the real experts — users — to choose the outstanding suppliers of network products and services. Don't miss this opportunity to cast your vote!

Thanks,

John Gallant

John Gallant
Editor-in-Chief

NETWORK WORLD
THE NEWSWEEKLY OF ENTERPRISE NETWORK COMPUTING

The Meadows, 161 Worcester Road, Framingham, MA 01701-9172, (508) 875-6400

An IDG Company: The World's Leader In Information Services Technology.

Same-day service. Anywhere in the country.
For three years. At no extra cost.

Mission-critical PC Server 85/95 service from IBM.* Getting it
from anyone else may be mission impossible.

PC Server
Model 95

IBM ®

Each week, *Network World* brings you news and views on the technologies shaping enterprise networks. Now it's your turn to speak out. NW's Enterprise Technology Awards give you the chance to spotlight suppliers whose products work best in a multi-vendor network environment.

This entry form lists major suppliers of LAN, WAN, software and other products. In each category, please select the ONE vendor whose product you feel is easiest to install,

use and manage in a multivendor network environment. A space marked "Other" is available for write-in candidates. The winning companies will be honored during a special ceremony at NetWorld + Interop in Atlanta in September.

This form is a prepaid mailer that can be returned at no expense to you. **ALL ENTRIES MUST BE RECEIVED BY JULY 1, 1994.**

No entry will be accepted without your name and company name.

---

1. Adapters
A. Ethernet
1❏ Accton Technology Corp. 05
2❏ Artisoft, Inc.
3❏ Asante Technologies, Inc.
4❏ Cabletron Systems, Inc.
5❏ Digital Communications Associates, Inc.
6❏ Digital Equipment Corp.
7❏ D-Link Systems, Inc.
8❏ Gateway Communications, Inc.
9❏ Hayes Mirocomputer Products, Inc.
0❏ IBM
x❏ Intel Corp.
y❏ NetWorth, Inc.
1❏ Proteon, Inc. 06
2❏ Racal-Datacom, Inc.
3❏ Standard Microsystems Corp.
4❏ Thomas-Conrad Corp.
5❏ 3Com Corp.
6❏ Tiara Computer Systems, Inc.
7❏ Other _____ 07

B. Token Ring
1❏ Andrew Corp. 08
2❏ Asante Technologies, Inc.
3❏ Digital Communications Associates, Inc.
4❏ IBM
5❏ Madge Networks, Inc.
6❏ NetWorth, Inc.
7❏ Olicom USA, Inc.
8❏ Optical Data Systems, Inc.
9❏ Proteon, Inc.
0❏ Racal-Datacom, Inc.
x❏ RAD Data Communications, Inc.
y❏ Standard Microsystems Corp.
1❏ Thomas-Conrad Corp. 09
2❏ 3Com Corp.
3❏ Other _____ 10

C. FDDI
1❏ Cabletron Systems, Inc. 11
2❏ Codenoll Technology Corp.
3❏ Cisco/Crescendo Communications, Inc.
4❏ Digital Equipment Corp.
5❏ Fibronics International, Inc.
6❏ IBM
7❏ Interphase Corp.
8❏ Network Peripherals, Inc.
9❏ Network Systems Corp.
0❏ Proteon, Inc.
x❏ SBE, Inc.
y❏ Summit Microsystems Corp.
1❏ SysKonnect, Inc. 12
2❏ 3Com Corp.
3❏ Other _____ 13

2. Bridges
1❏ Advanced Computer Communications 14
2❏ Andrew Corp.
3❏ Cabletron Systems, Inc.
4❏ CrossComm Corp.
5❏ Digital Equipment Corp.
6❏ Hewlett-Packard Co.
7❏ IBM
8❏ Magnalink Communications Corp.
9❏ Microcom, Inc.
0❏ Netronix, Inc.
x❏ Network Systems Corp.
y❏ Newbridge Networks, Inc.
1❏ Olicom USA, Inc. 15
2❏ Plexcom, Inc.
3❏ RAD Data Communications, Inc. (continued)

4❏ RAD Network Devices Inc.
5❏ Retix
6❏ 3Com Corp.
7❏ Ungermann-Bass, Inc.
8❏ Xyplex, Inc.
9❏ Other _____ 16

3. Routers
1❏ Advanced Computer Communications 17
2❏ Ascom Timeplex, Inc.
3❏ Cisco Systems, Inc.
4❏ CrossComm Corp.
5❏ Digital Equipment Corp.
6❏ Hypercom, Inc.
7❏ IBM
8❏ NEC America, Inc.
9❏ Network Systems Corp.
0❏ Novell, Inc.
x❏ Plexcom, Inc.
y❏ Proteon, Inc.
1❏ RAD Network Devices, Inc. 18
2❏ Retix
3❏ 3Com Corp.
4❏ Wellfleet Communications, Inc.
5❏ Xyplex, Inc.
6❏ Other _____ 19

4. Intelligent Hubs
1❏ Alantec 20
2❏ Artel Communications Corp.
3❏ Asante Technologies, Inc.
4❏ Ascom Timplex, Inc.
5❏ Cabletron Systems, Inc.
6❏ Chipcom Corp.
7❏ David Systems, Inc.
8❏ Fibronics International, Inc.
9❏ Fibermux Corp.
0❏ Gandalf Systems Corp.
x❏ Hewlett-Packard Co.
y❏ Hughes LAN Systems, Inc.
1❏ IBM 21
2❏ Klever Computers
3❏ LANNET Data Communications, Inc.
4❏ NCR Corp.
5❏ NetWorth, Inc.
6❏ Optical Data Systems, Inc.
7❏ Plexcom, Inc.
8❏ Proteon, Inc.
9❏ Racal-Datacom, Inc.
0❏ RAD Data Communications, Inc.
x❏ Standard Microsystems Corp.
y❏ SynOptics Communications, Inc.
1❏ 3Com Corp. 22
2❏ Ungermann-Bass, Inc.
3❏ Xyplex, Inc.
4❏ Other _____ 23

5. LAN Diagnostic/Management Tools
1❏ BICC Data Networks, Inc./3Com Corp. 24
2❏ Brightwork Development
3❏ Data General Corp.
4❏ Digilog, Inc.
5❏ Frye Computer Systems, Inc.
6❏ Hewlett-Packard Co.
7❏ IBM
8❏ Legent Corp.
9❏ Network Communications Corp.
0❏ Network General Corp.
x❏ NetWorth, Inc.
y❏ Novell, Inc.
1❏ ProTools 25
2❏ Scope Communications (continued)

3❏ Spider Systems, Inc.
4❏ Other _____ 26

6. LAN Servers
1❏ Advanced Logic Research, Inc. 27
2❏ AST Research, Inc.
3❏ Compaq Computer Corp.
4❏ CompuAdd Corp.
5❏ Dell Computer Corp.
6❏ Digital Equipment Corp.
7❏ Everex Systems, Inc.
8❏ Hewlett-Packard Co.
9❏ IBM
0❏ Network Connections, Inc.
x❏ Northgate Computer Corp.
1❏ Zenith Data Systems Corp. 28
2❏ Other _____ 29

7. Superservers
1❏ Advanced Logic Research, Inc. 30
2❏ Acer, Inc.
3❏ Altos Computer Systems
4❏ Auspex Systems, Inc.
5❏ Compaq Computer Corp.
6❏ Digital Equipment Corp.
7❏ Epoch Systems, Inc.
8❏ Hewlett-Packard Co.
9❏ IBM
0❏ NCR Corp.
x❏ NetFRAME Systems, Inc.
y❏ Parallan Computer, Inc.
1❏ Tricord Systems, Inc. 31
2❏ Unisys Corp.
3❏ Other _____ 32

8. Network Operating Systems
1❏ Banyan Systems, Inc. 33
2❏ Digital Equipment Corp.
3❏ IBM
4❏ Microsoft Corp.
5❏ Novell, Inc.
6❏ Other _____ 34

9. Peer-to-Peer Networks
1❏ Apple Computer, Inc. 35
2❏ Artisoft, Inc.
3❏ Gateway Communications, Inc.
4❏ Hayes Microcomputer Products, Inc.
5❏ IBM
6❏ Invisible Software, Inc.
7❏ Microsoft Corp.
8❏ Novell, Inc.
9❏ Sitka Corp.
0❏ Tiara Computer Systems, Inc.
x❏ Webcorp
y❏ Other _____ 36

10. Wireless LANs
1❏ AT Schindler 37
2❏ BICC Communications/3COM Corp.
3❏ Caliber Tek, Inc.
4❏ California Microwave
5❏ INFRALINK of America, Inc.
6❏ Motorola, Inc.
7❏ NCR Corp.
8❏ O'Neill Communications, Inc.
9❏ Photonics Corp.
0❏ Proxim, Inc.
x❏ Telesystems SLW, Inc.
y❏ Windata, Inc.
1❏ Other _____ 38

11. Apple Connectivity Products
1❏ Asante Technologies, Inc. 39
2❏ Alisa Systems, Inc.
3❏ Apple Computer, Inc.
4❏ Artisoft, Inc.
5❏ Cayman Systems, Inc.
6❏ Compatible Systems, Inc.
7❏ Dayna Communications, Inc.
8❏ Digital Equipment Corp.
9❏ Farallon Computing, Inc.
0❏ Helios USA
x❏ NEC America, Inc.
y❏ Novell, Inc. (continued)

1❏ Pacer Software, Inc. 40
2❏ Shiva Corp.
3❏ Sitka Corp.
4❏ Webster Computer Corp.
5❏ Other _____ 41

12. Remote LAN Access
1❏ Cayman Systems, Inc. 42
2❏ Centrum Communications, Inc.
3❏ Citrix Systems, Inc.
4❏ Combinet, Inc.
5❏ DigiBoard
6❏ Digital Communications Associates, Inc.
7❏ IBM
8❏ Microcom, Inc.
9❏ Multi-Tech Systems, Inc.
0❏ NEC America, Inc.
x❏ Norton-Lambert Corp.
y❏ Novell, Inc.
1❏ RAD Data Communications, Inc. 43
2❏ Rockwell International Corp.
3❏ Shiva Corp.
4❏ Symantec Corp.
5❏ Telebit Corp.
6❏ U.S. Robotics
7❏ Xircom, Inc.
8❏ Xylogics, Inc.
9❏ Other _____ 44

13. Imaging Systems
1❏ Compulink Management Center, Inc. 45
2❏ Digital Equipment Corp.
3❏ FileNet Corp.
4❏ Formtek, Inc.
5❏ IBM
6❏ Imnet, Inc.
7❏ Indus, Int'l
8❏ Kodak
9❏ Optika Imaging Systems, Inc.
0❏ Recognition International, Inc.
x❏ Simplify Development Corp.
y❏ Sigma Systems, Inc.
1❏ View Star Corp. 46
2❏ Wang Laboratories, Inc.
3❏ Other _____ 47

14. Database Servers
1❏ Borland International, Inc. 48
2❏ Cincom Systems, Inc.
3❏ Computer Associates International, Inc.
4❏ Computer Corp. of America
5❏ Digital Equipment Corp.
6❏ Gupta Technologies, Inc.
7❏ Hewlett-Packard Co.
8❏ IBM
9❏ Information Builders, Inc.
0❏ Informix Software, Inc.
x❏ Ingres Corp.
y❏ Micro Data Base Systems, Inc.
1❏ Microsoft Corp. 49
2❏ Novell, Inc.
3❏ Oracle Corp.
4❏ Progress Software Corp.
5❏ Revelation Technologies, Inc.
6❏ Software AG
7❏ Sybase, Inc.
8❏ Unify Corp.
9❏ XDB Systems, Inc.
0❏ Other _____ 50

15. E-Mail Software
1❏ Banyan Systems, Inc. 51
2❏ Beyond, Inc.
3❏ CE Software, Inc.
4❏ Data General Corp.
5❏ DaVinci Systems Corp.
6❏ Digital Equipment Corp.
7❏ Enable Software
8❏ Futurus Corp.
9❏ Hewlett-Packard Co.
0❏ IBM
x❏ Lotus Development Corp./cc:Mail
y❏ Microsoft Corp.
1❏ Notework Corp. 52
2❏ Reach Software Corp. (continued)

3❏ Other _____ 53

16. Group Collaboration Software
1❏ Campbell Services, Inc.   54
2❏ Collaborative Technologies Corp.
3❏ Coordination Technology, Inc.
4❏ Digital Equipment Corp.
5❏ Enable Software
6❏ Forefront Group
7❏ Futurus Corp.
8❏ Group Technologies, Inc.
9❏ Lotus Development Corp.
0❏ Microsoft Corp.
x❏ Micro Systems Software
y❏ NCR Corp.
1❏ ON Technology, Inc.   55
2❏ Powercore, Inc.
3❏ Ventana Corp.
4❏ WordPerfect Corp.
5❏ Other _____ 56

17. Workflow
1❏ Action Technologies, Inc.   57
2❏ Delrina Corp.
3❏ Filenet Corp.
4❏ IBM
5❏ Reach Software Corp.
6❏ Other _____ 58

18. Client/Server Applications Development Tools
1❏ Anderson Consulting   59
2❏ Blyth Software, Inc.
3❏ Cadre Technologies, Inc.
4❏ Cognos Corp.
5❏ Gupta Corp.
6❏ IBM
7❏ Inference Corp.
8❏ Informix Software, Inc.
9❏ Ingres Corp.
0❏ Object Systems
x❏ Oracle Corp.
y❏ Powersoft Corp.
1❏ Progress Software Corp.   60
2❏ Seer Technologies, Inc.
3❏ Spinnaker Software Corp.
4❏ Sybase, Inc.
5❏ Texas Instruments, Inc.
6❏ Uniface Corp.
7❏ Visix Software, Inc.
8❏ Other _____ 61

19. Integrated Network Management Systems
1❏ Applied Computing Devices, Inc.   62
2❏ Boole & Babbage, Inc. (continued)
3❏ Cabletron Systems, Inc.
4❏ Digital Equipment Corp.
5❏ Hewlett-Packard Co.
6❏ IBM
7❏ MAXM Systems Corp.
8❏ NetLabs, Inc.
9❏ NCR Corp.
0❏ Novell, Inc.
x❏ Nynex Allink Co.
y❏ Objective Systems Integrators
1❏ Sprint Corp.   63
2❏ SunConnect
3❏ Systems Center
4❏ Other _____ 64

20. DSU/CSU
1❏ ADC Kentrox (continued)   65

2❏ ADTRAN
3❏ AT&T Paradyne
4❏ Data Race, Inc.
5❏ Digital Access Corp.
6❏ Digital Link Corp.
7❏ Dowty Communications, Inc.
8❏ General DataComm, Inc.
9❏ Larscom Inc.
0❏ Motorola Inc. (Codex, UDS)
x❏ Penril DataComm Networks
y❏ Racal-Datacom, Inc.
1❏ RAD Data Communications, Inc.   66
2❏ Teleprocessing Products
3❏ TxPORT
4❏ Verilink Corp.
5❏ Other _____ 67

21. High-Speed Modems (9.6K & above)
1❏ AT&T Paradyne   68
2❏ Canoga-Perkins
3❏ General DataComm, Inc.
4❏ Hayes Microcomputer Products, Inc.
5❏ Microcom, Inc.
6❏ Motorola Inc. (Codex, UDS)
7❏ Multi-Tech Systems, Inc.
8❏ Octocom Systems, Inc.
9❏ Plexcom, Inc.
0❏ Practical Peripherals, Inc.
x❏ Racal-Datacom, Inc.
y❏ Telebit Corp.
1❏ Zypcom, Inc.   69
2❏ US Robotics, Inc.
3❏ Zyxel, Inc.
4❏ Other _____ 70

22. Communications Software
1❏ Datastorm Technologies, Inc.   71
2❏ Digital Communications Associates, Inc.
3❏ Farallon Computing, Inc.
4❏ FTP Software, Inc.
5❏ Microcom, Inc.
6❏ Microsoft Corp.
7❏ Network Management, Inc.
8❏ Norton Lambert Corp.
9❏ Software Ventures Corp.
0❏ Symantec Corp.
x❏ Triton Technologies, Inc.
y❏ Walker Richer & Quinn, Inc.
1❏ Wollongong Group, Inc.   72
2❏ Other _____ 73

23. Multiplexers
1❏ Ascom Timplex, Inc.   74
2❏ AT&T Paradyne (continued)
3❏ CASE Communications, Inc.
4❏ Data Race, Inc.
5❏ Distinct Corp.
6❏ Frontier Technology
7❏ Gandalf Systems Corp./Infotron
8❏ General DataComm, Inc.
9❏ Micom Communications Corp.
0❏ Motorola, Inc. (Codex)
x❏ Multi-Tech Systems, Inc.
y❏ Netrix Corp.
1❏ Network Equipment Technologies, Inc.   75
2❏ Newbridge Networks, Inc.
3❏ Plexcom, Inc.
4❏ Racal-Datacom, Inc.
5❏ RAD Data Communications, Inc.
6❏ StrataCom, Inc.
7❏ Telco Systems, Inc. (continued)

8❏ Tellabs Operations, Inc.
9❏ T3 Plus Networking, Inc.
0❏ Other _____ 76

24. Packet Switches
1❏ Alcatel Business Systems, Inc.   77
2❏ Ascom Timplex, Inc.
3❏ BBN Communications Corp.
4❏ BT N.A.
5❏ Dynatech Communications, Inc.
6❏ Hughes Network Systems, Inc.
7❏ Motorola Inc. (Codex)
8❏ Netrix Corp.
9❏ Northern Telecom, Inc.
0❏ Plantronics, Inc.
x❏ RAD Data Communications, Inc.
y❏ Sprint International
1❏ Telematics International, Inc.   78
2❏ Tellabs Operations, Inc.
3❏ Other _____ 79
                         80-1

25. PBXs
1❏ AT&T   205
2❏ Ericsson Information Systems
3❏ Fujitsu Business Communications Systems, Inc.
4❏ Harris Corp.
5❏ Hitachi America, Ltd.
6❏ InteCom, Inc.
7❏ Mitel, Inc.
8❏ NEC America, Inc.
9❏ Northern Telecom, Inc.
0❏ Rolm
x❏ Siemens Private Communications Systems
y❏ Other _____ 06

26. Videoconferencing Systems
1❏ AT&T   07
2❏ BT N.A.
3❏ Compression Labs, Inc.
4❏ GPT Video Systems
5❏ Grass Valley Group, Inc.
6❏ Harris Corp.
7❏ Hitachi America, Ltd.
8❏ Intel Corp.
9❏ NEC America, Inc.
0❏ Oki America, Inc.
x❏ Panasonic Corp.
y❏ PictureTel Corp.
1❏ Videoconferencing Systems, Inc.   08
2❏ VideoTelecom Corp.
3❏ Vtel Corp.
4❏ Other _____ 09

27. Voice Messaging Systems
1❏ Active Voice Corp.   10
2❏ Applied Voice Technology, Inc.
3❏ AT&T
4❏ Boston Technology, Inc.
5❏ Centigram Communications Corp.
6❏ Digital Sound Corp.
7❏ Microlog Corp.
8❏ Northern Telecom, Inc.
9❏ Octel Communications Corp.
0❏ Rolm Co.
x❏ VMX, Inc.
y❏ Other _____ 11

28. Digital Private Line Services
1❏ Advanced Telecommunications Corp.   12
2❏ American Private Line Services, Inc.
3❏ AT&T (continued)

4❏ Cable & Wireless Communications, Inc.
5❏ Consolidated Network, Inc.
6❏ MCI Communications Corp.
7❏ Metromedia Communications Corp.
8❏ Sprint Corp.
9❏ WilTel
0❏ Other _____ 13

29. 800 Services
1❏ Advanced Telecommunications Corp.   14
2❏ Allnet Communications Services, Inc.
3❏ AT&T
4❏ Cable & Wireless Communications, Inc.
5❏ LCI International
6❏ MCI Communications Corp.
7❏ Metromedia Communications Corp.
8❏ RCI Long Distance
9❏ Sprint Corp.
0❏ WilTel
x❏ Other _____ 15

30. Frame Relay Services
1❏ AT&T   16
2❏ BT N.A.
3❏ Cable & Wireless Communications, Inc.
4❏ CompuServe, Inc.
5❏ Graphnet, Inc.
6❏ GTE Telephone Operations
7❏ MCI Communications Corp.
8❏ Sprint Corp.
9❏ Wiltel
0❏ Other _____ 17

31. Switched Digital Services
1❏ Allnet   18
2❏ AT&T
3❏ Cable & Wireless Communications, Inc.
4❏ MCI Communications Corp.
5❏ Sprint Corp.
6❏ Wiltel
7❏ Other _____ 19

32. Value-Added Network Services
1❏ Advantage Systems, Inc.   20
2❏ Advantis (IBM Information Network/Sears Communications Co.)
3❏ AT&T
4❏ BT N.A.
5❏ CompuServe, Inc.
6❏ GE Information Services
7❏ Graphnet, Inc.
8❏ Infonet Computer Sciences Corp.
9❏ MCI Communications Corp.
0❏ Sprint Corp.
x❏ Other _____ 21

33. Virtual Network Services
1❏ AT&T   22
2❏ Cable & Wireless Communications, Inc.
3❏ Infonet Computer Sciences Corp.
4❏ MCI Communications Corp.
5❏ Sprint Corp.
6❏ Other _____ 23

NAME:_____

TITLE:_____

COMPANY:_____

PHONE:_____

ADDRESS:_____